Journal of Computer Security 10 (2002) 1–3
IOS Press

Guest editor's preface

In November of 2000, it was my privilege to serve as Workshop Chair for the First Workshop on Intrusion Detection and Prevention, held under the auspices of the 7th ACM Conference on Computer Security. The experience proved to be a memorable one for several reasons, including the wonderful hospitality of the organizers, and the rich history of the surrounding area. This setting of combined ancient architectures and modern ideas seemed a particularly appropriate site for our workshop; in devising intrusion detection systems; we rely simultaneously on time-honored strategies for protection, identification and defense, and on employing the latest hardware and software that can be obtained!

From among those speakers who participated, we selected the following papers for expansion. The authors responded with the results you see here. The selected subjects also reflect the contrast of "old and new" relative to the development of the field of IDS – for instance, we have papers involving of profiling, a tried-and-true strategy for identifying potential misuse, as well as a discussion of the relatively recent "business model" of security. A brief description of the subject matter is provided below. I hope that you will benefit as much from reading them, as I have from working with their authors.

As IDS become mainstream, businesses increasingly find that previous ad hoc methods for determining "how much" computer security is "enough" are insufficient. The paper "Toward cost-sensitive modeling for intrusion detection and response", by W. Lee, W. Fan, M. Miller, S.J. Stolfo and E. Zadok, seeks to put the issue of cost-benefit analysis of IDS on a formal basis. Their work combines a range of costs – from development to operational, and successful and unsuccessful diagnosis of attacks – into a single cost model. This paper additionally includes results of empirical experiments, and preliminary results indicate that use of this model can be effective in reducing overall "costs" of intrusion detection.

"Using internal sensors and embedded detectors for intrusion detection", by F. Kerschbaum, E.H. Spafford and D. Zamboni, adds to our understanding of the utility of internal sensors as means of detecting intrusions (and misuse) in computer systems. This paper provides several useful results. One is the classification of data collection strategies for intrusion detection systems, and a discussion of the strengths and weaknesses of each. Another is the introduction of a framework for building IDS based on internal sensors – the ESP architecture, and a discussion of a prototype implementation. This paper also includes analysis of both the effectiveness of the technique (high for those attacks it was designed to detect) and the performance costs associated with using embedded detectors.

Another paper, "STATL: An attack language for state-based intrusion detection", by S.T. Eckmann, G. Vigna and R.A. Kemmerer at UC Santa Barbara, adds to our understanding of the state transition analysis technique piloted in 1992 by Dr. Kemmerer and his students. Positive results from the analysis of the earlier prototypes of the concept and recognition of the similarities among them has led to the development of a domain-independent language – STATL – as described in this paper. STATL has a precise syntax and parser, and a formally defined semantics, and is extensible. The STATL work should be of interest both to those interested in using STATL to support their own intrusion detection work, and for those who seek a way to combine IDS or translate between IDS.

Detection systems often focus on their ability to detect "short duration" attacks – those that take place within a limited period of time. There is increasing interest in detecting stealthy attacks, and stealthy preludes to attacks – sometimes referenced as "low and slow" activities. Detection strategies for these attacks tend to be plagued by high costs of data maintenance and large numbers of false positives. The paper by S. Staniford, J.A. Hoagland and J.M. McAlerney, "Practical automated detection of stealthy portscans", discusses a method whereby stealthy portscans can be detected efficiently. This method is implemented in the SPICE (Stealthy Probing and Intrusion Correlation Engine) portscan detector. The basis of the author's technique is to use simulated annealing to form clusters of packets with similar "anomalousness" estimations, and retain those with high "anomalous" ratings for a longer period of time than those that appear normal. Results indicate that this strategy can detect a wide range of scans, including stealthy scans, without unacceptably high false positive rates.

Profiling is one of the more frequent strategies utilized by IDS developers, being a useful tool for wide range of detection activities. However, despite the frequency of use, the concept is still an active area of research. The paper "Enhancing profiles for anomaly detection using time granularities", by Y. Li, N. Wu, X.S. Wang and S. Jajodia of George Mason University, expands our understanding of how "time" affects profiles. Time-dependency (particularly knowledge about work schedules) has long been known to affect behavior profiles, but usually this information is utilized in an ad hoc way. This paper introduces the notion of calendar schema for describing typical time interval patterns, and temporal association rules for combining traditional association rules with the calendar based patterns. Also discussed is an extension of the Apriori algorithm for building temporal profiles, and some early experimental results based on applying the authors' strategy to the 1998 DARPA intrusion detection evaluation data.

D. Spinellis and D. Gritzalis, in "Panoptis: Intrusion detection using a domain-specific language", provides an example of how one might use a specialized language to express the design values and constraints of an IDS. Their application, Panoptis, uses Unix process-accounting records as the basis for an anomaly detection system that incorporates profiling of execution entities; the authors use a domain-specific language to indicate the particular aspects of execution that will be checked. The

architecture proposed in this paper is intended to be useful for allowing focused detection systems such as Panoptis to work together in a confederation to detect misuse of many kinds, including wiretapping, information leakage, tampering, and masquerading.

"An environment for security protocol intrusion detection", written by A. Yasinsac, describes an approach to detecting attacks on security protocols in real time. The author discusses a tool set architecture for analyzing a security protocol. The approach advocated in this paper utilizes both recognition of the characteristics of specific known attacks (including known classes of attacks), and recognition of activity that is inherently suspect. Attack characterization is based on sequences of activity; patterns of protocol sequences are used to devise attack signatures in a similar way that code patterns in files may be used to devise virus signatures.

C.R. Ramakrishnan and R. Sekar, in "Model-based analysis of configuration vulnerabilities", provide new insight into the age-old problem of how best to perform automated computer system configuration analysis. Earlier research is best typified by the approaches used in COPS and SATAN, both of which assist the user in determining whether known exploitable conditions exist. In contrast, the approach here uses recent improvements in model-checking strategies; it compares the formal specification of desirable security properties with an abstract model of a system's security-related behaviors, and determines whether the model satisfies the properties. An examplar system, a simplified version of Unix, is used to illustrate the effectiveness of the technique.

Deborah Frincke
Center for Secure and Dependable Systems
University of Idaho