CERIAS Tech Report 2001-97
New directions for the AAFID architecture
by Eugene H. Spafford and Diego Zamboni
Center for Education and Research
Information Assurance and Security
Purdue University, West Lafayette, IN 47907-2086

New directions for the AAFID architecture

Topic category:
"Innovative Approaches: New results related to innovative ways of thinking about IDS"

Eugene Spafford, Diego Zamboni
Center for Education and Research in Information Assurance and Security
Purdue University
West Lafayette, IN 47907-1398, U. S. A.
{spaf,zamboni}@cs.purdue.edu
Ph. (765) 494-7805, Fax (765) 494-0739

Talk proposal submitted to the
Recent Advances in Intrusion Detection 99 Workshop
Time requested: 30 minutes

Author biographies:

Eugene Spafford

Gene Spafford is a professor of Computer Sciences at Purdue University. He is the founder and director of the COAST Laboratory, and the new Center for Education and Research in Information Assurance and Security (CERIAS). He has been involved in Intrusion Detection and avoidance research for over a decade.

Diego Zamboni

Diego Zamboni is a Ph.D. student at Purdue University, where he is working in CERIAS in Intrusion Detection research. He obtained his M.S. in Computer Science from Purdue University. Previously he obtained his bachelor's degree in Computer Engineering from the National Autonomous University of Mexico, where he also established one of the first Computer Security Incident Response Teams in Mexico.

New directions for the AAFID architecture

Topic category suggested:
"Innovative Approaches: New results related to innovative ways of thinking about IDS"

At RAID'98 we presented the Autonomous Agents for Intrusion Detection (AAFID) architecture and prototype. Since that time, experience with the prototype has allowed us to identify the following major problems:

Agent generality: The original AAFID prototype implemented a polling-loop structure for the agents, which covers many possible applications, but is insufficient for more complex monitoring tasks where, for example, the agent needs to block on a certain system resource or wait for a specific signal to occur.

Performance: Deploying the AAFID prototype in a production environment has shown that the current Perl implementation requires considerable system overhead, easily overloading the host when more than a few agents are running.

Data analysis: The original AAFID prototype is a data-collection engine, but has no ability to interpret the data it receives, which is simply collected at the top-level monitor. For an IDS to be truly useful, some form of analysis of the data should be performed. However, many of the conventional techniques used for analysis in Intrusion Detection Systems cannot be directly applied when we have data coming from many different agents on many different hosts.

We have been working on the problems mentioned, resulting in a number of developments that will be described in this talk.

First, a number of improvements have been made to the AAFID prototype, making agents more powerful and easy to write. The original polling-loop structure of the agents has been dropped in favor of a more general event-handling structure, which allows the agents to utilize fewer system resources and look for a wider range of events. We have also been working on porting the prototype to Windows NT, and we will report on the results of that effort.

Second, we are working on improving the current AAFID implementation to make it less resource-intensive. Some lines of work in this respect include making each agent a thread instead of a separate process, and allowing the implementation of agents in languages other than Perl.

Finally, we have started studying different approaches to analyzing the data coming from a large number of distributed sensors. Some of these techniques include machine learning techniques such as clustering, and preliminary results are encouraging, suggesting that for some types of data and for some attacks, we may be able to automatically classify the data coming from the agents as "good" and "bad" without having to manually label or correlate each detected event.

These developments constitute important steps in the development of AAFID, both as a research prototype for innovative intrusion detection techniques and as a practical tool for intrusion detection. We will describe our work and findings in all the areas described, as well as paint a picture of the work that still lies ahead.