Data Mining Support for Intrusion Detection and Prevention

Richard Wasniowski

Computer Science Department
California State University Dominguez Hills
Carson, CA 90747

Abstract

As computer attacks are becoming more and more difficult to identify the need for better and more efficient intrusion detection systems increases. The main problem with current intrusion detection systems is high rate of false alarms. In this paper we discuss our experience in analyzing potential benefits of distributed multi sensor systems for intrusion detection. Our purpose for this work is to examine how to integrate multiple intrusion detection sensors in the order to minimize the number of incorrect-alarms

1. Introduction

The main problem with current intrusion detection systems is high rate of false alarms triggered off by attackers. Effective protecting the network against malicious attacks remains problem in both research and the computer network managing professionals. Improved monitoring of malicious attacks will require integration of multiple monitoring systems. In our current project we are analyzing potential benefits of distributed multi sensor systems for intrusion detection. The first problem is how to integrate data from multiple sensors, and the second how to identify most important data provided by multiple sensors. We are developing series of analytical and mathematical models to use potential benefits of multiple sensors for reducing false alarms.

The purpose of this paper is to discuss implementation of prototype multi sensor based intrusion detection system. We are especially interested in analyzing traffic that has an abnormal or malicious character and should prompt a closer look. A specific feature of the model is that the systems use multiple sensors and mining to process log files. This reduces the overhead in a distributed intrusion detection system.

2. Background

There are several intrusion detection systems, and one of the most popular in public domain is Snort[4]. Snort looks for attack signatures, which are specific patterns of activity that has been defined to be of a suspicious or malicious intent. Snort relies on the ability to recognize attack signatures in order to identify an attack. These pattern recognition definitions are called rules. Attacks are not static, as they are continuously evolving as systems are protected to withstand existing attack methodologies. As indicated by Cox and Gerg [7], Snort is an open source network packet monitoring and Intrusion Detection System. Snort looks for attack signatures, which are specific patterns of activity that has been defined to be of a suspicious or malicious intent. Snort analyzes network packets, and thus is classified as a Network Intrusion Detection System, or NIDS. These types of systems must be connected to the networks that they monitor and unless the network topology is very simple, multiple Snort systems, called Snort sensors, must be setup and configured to monitor these networks. Snort relies on the ability to recognize attack signatures in order to identify an attack. These pattern recognition definitions are called rules. Attacks are not static, as they are continuously evolving as systems are protected to withstand existing attack methodologies. Thus, it is critical to perform analysis of prior activity to look for trends or changes in activity that are not typically classified as an attack which are often the precursor to an attack. Though it is possible to analyze information on multiple Snort sensors one at a time, it is difficult to summarize or perform analysis from a multi-Snort sensor perspective. For example, if an organization has multiple office locations in widely different geographical locations, it would be expected that separate Snort sensors are configured and operating.

If an attacker targets the organization, it is possible that these different geographical locations are probed and

attacked in series or simultaneously. Being able to recognize probing or attacking at a multi-geographic perspective can provide value in understanding individual sensor alerts. Snort evaluates data at the packet level and thus must process a large amount of data in real-time. Because of this, logging performance is important and to achieve this, the data storage implementation for Snort is optimized for fast writing. This results in a highly normalized database design where there are many one to one relationships. In fact, information representing the primary type of information in Snort, called an event (which is the packet information that matched one or more Snort rules), is represented in no less than six tables. One of the tables contains information that must always be provided for every event, and the other five tables contain information that is optional depending on the type of event that has occurred. This implementation allows for writing the smallest amount of information at a time, which allows for high performance when logging information. However, this design's drawback is when there is a need to read the information for reporting and other analysis. Displaying information for a single event may require joining six or more tables using outer joins, which impacts reporting performance. The primary reason for building a data mart or data warehouse is to develop an intelligent, consolidated view of enterprise information. But each year, a large number of business intelligence and data warehousing initiatives fail because of erroneous or incomplete data. Often, users ignore the importance of developing a data management strategy as part of their extract, transform, and load or data warehouse architecture. Even with this highly normalized database design, the log data cannot be kept indefinitely, requiring that the data is removed from the sensor system by deleting older data. This results in the loss of data that could be used to develop better rules or provide evidence of an attack. Though it can be archived before deleting, the data is then offline and harder to analyze. Existing multi-Snort log reporting applications do exist. ACID, a popular web-database application has been available since. However, ACID is designed so that it can be configured as the primary store for Snort log data and thus is subject to the same performance and historical data issues that the Snort sensors face - in the section titled "The Ongoing Use of the ACID Console," Cox and Gerg [7] discusses deleting Snort log data on a periodic basis, though recommending backing up the data to some type of offline storage before deleting the data.

A honeypot is a system whose value lies in being probed, attacked, or otherwise taken advantage of by attacker. Spitzner classifies honeypot solutions into two broad categories: production and research. For research purposes, we simply want to collect as much information on our attackers as possible. Production systems are generally used as an added layer of network security. Production honeypots are thought of as simpler and more intuitive than research honeypots. This affords system administrators the freedom to select from several commercially available (and sometimes free) honeypot solutions. Examples of such solutions that are currently available include BackOfficerFriendly, Specter, and honeyd. Research honeypots, on the other hand, are often homemade solutions that can track an attacker's actions. Network security professionals and educational institutions often employ research honeypots in the hopes of seeing a hacker in action. A honeypot that is to be used for research will often contain a fully operational operating system running certain services and vulnerabilities. Generally, this type of honeypot is much more difficult to configure and requires more time for upkeep. AIDE is a tool for trying to detect if someone has been on our machine and changed anything. If we know or suspect that someone has been on our machine, we can run aide to see what files have been modified, this will be a great help as we try to see what someone has done. AIDE works be creating a checksum of files in specific, user defined directories, saving these in a database, and checking them against the same files at a later date. The major drawback to AIDE is human, that is, it has to be run BEFORE an attack or it is worthless.

3. System design

Intrusion detection monitoring multiple systems and networks requires the existence of multiple intrusion detection systems. Each network being monitored requires its own intrusion detection system. Also, bandwidth limitations more than one intrusion detection system may be required on the same network.

In practical functioning intrusion detection systems produce false alarms called "false positive." This is basically an event identified as an intrusion attempt, but in reality it is not. The typical response to this by the administrator is to reconfigure the intrusion detection system to not identify that particular event as an intrusion attempt.

On the other hand, being constantly notified by "false positives" may also result in a false sense of security, as the administrator can adopt an attitude that typically intrusion events reported by the intrusion detection system are false positives and may not properly respond

to a real intrusion attempt. In order to be more responsive the intrusion detection system must be configured in a way that will probably report many of these "false positives".

Additional tools are required to better understand all of the data generated by the intrusion detection systems. This need is not new, there are many existing solutions that will read from and produce reports from intrusion detection system logs [see discussion on various forums [19-24] One possibility to minimize 'false positives" is to fuse data from multiple sensors. This requires both new data fusion methods and practical experiments. The main goal of our work is to develop such methods and test them in experimental setup. This report describes first step in designing and implementing such a system.
The snort system uses various rule-based techniques based on comparison of past and current attacks. In order to implement efficient intrusion detection system integration of multiple techniques and tools with snort is required. The real time data collection process is very intensive and produces large amount of data. In addition the whole system depends on data and database failure must be prevented. "Recognizing whether two sensors see the same or two different objects is a major challenge. Another challenge is effectively processing data streams that come from multiple sensors. Features that are not typical of the traditional database management system, such as almost real-time response. The general framework for the system is shown on Fig.1, and placing sensors in three different configurations is shown on Fig.2-3.
Fusion model similar to fwsnort [18] translates snort rules into an equivalent iptables ruleset.

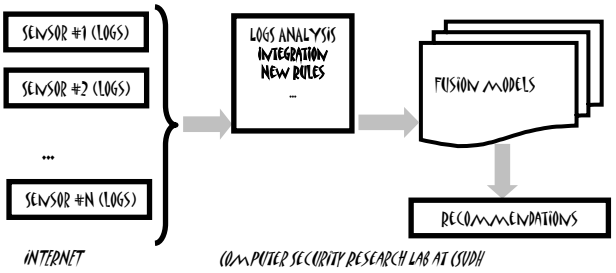

Fig. 1 General framework

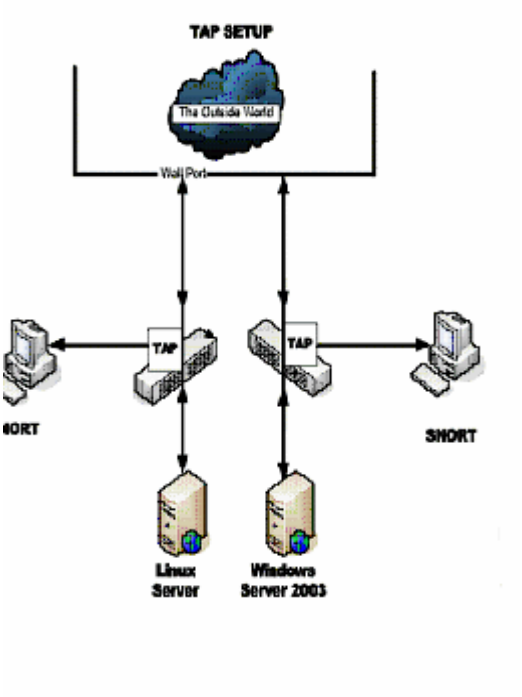

Fig. 2 Intrusion detection using taps

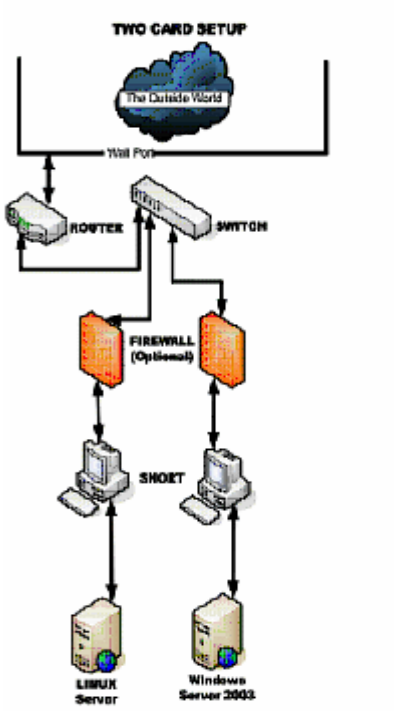

Fig. 2 Intrusion detection using bridge

4. Implementation

The Snort based multiple sensors system monitors two networks. One fully operational and one experimental. The NAT Router/Firewall is configured to allow Internet access to the web servers, by mapping selected ports to the web server behind the NAT Router/Firewall on the internal network. This configuration is selected to allow a single attack to simultaneously attack the NAT Router/Firewall and the web server so we could generate Snort events that had identical timestamps to ensure that we could successfully merge data from multiple snort sensors with identical timestamps. One web server was an Intel-based PC running Microsoft Windows 2003 Server, the second Centos based Linux system. The attack system is an Intel-based PC running Fedora Core 4 (FC4) GNU/Linux, laptop computer. Nmap was selected as it was considered by Cox and Gerg [7], as one of the most widely used port scanners for network analysis. Each Snort sensor is an Intel-based PC running CENTOS4.3 with Snort 2.3.0/2.6.0 and mySQL 4.3.10. Each Snort sensor is configured with identical rule sets (the set of rules included with Snort 2.3.0), to run in Intrusion Detection System (IDS) mode, and to log to the MySQL database engine installed on each Snort sensor. As indicated by Beale et al [4], logging to a relational database was selected as it is considered to be more efficient than logging to files, and later logging to file was added as it is useful for analyzing data by some specialized packages such as Snortalog. The system is implemented using Open Software whenever possible such as Snort, mySQL etc. We are using Windows 2003 servers for our Web server and Honeypot. Our intrusion detection sensors are installed on Linux based systems.

5. Findings and future work

As computer attacks become more and more sophisticated, the need to provide effective intrusion detection methods increases. Network-based distributed attacks are especially difficult to detect and require coordination among different intrusion detection components. We propose a solution that responds to such requirements. Our implemented NIDS model is in fact a prototype and needs to evolve into more mature and efficient model. Future work should emphasize a revisit of database design. One of the key reasons that the entities have so many attributes, in current implementation, was the concern of including important attributes and thus having all data available. This resulted in the inclusion of practically all of the event data. We believe that a good approach for achieving this

would be an expansion of the solution: including and consolidated version of the operational Snort database that is it is used in conjunction with NIDS for reporting and analysis. On the whole, our information fusion based intrusion detection model is in fact a prototype and needs to evolve into more mature and efficient model. Future work emphasizes a revisit of database design to allow data fusion from multiple

This project is described in details on the following password protected web site:
csc09.csudh.edu/csl

Conference organizers will receive password access if requested.

We have collected a huge amount of data such as alert logs and multiple mySQL databases and improved snort rules design and we are currently finalizing processing those sets of data.

12. Acknowledgement

Research is partially supported by NGA and DoD.

13. References

[1] John Ashenfelter, Data Warehousing with MySQL, http://www.mysqluc.com/cs/mysqluc2005/vi
[2] S. Axelsson. "Intrusion Detection Systems: A Taxonomy and Survey." Technical Report No 99-15, Dept of Computer Engineering, Chalmers University of Technology, Sweden, March 2000
[3] D. Paul Benjamin, Ranjita Shankar-Iyer, Archana Perumal, VMSoar: A Cognitive Agent for Network Security, http://csis.pace.edu/robotlab/pubs/VMSoarS.pdf
[4] Beale, Jay, 2004. Snort 2.1 Intrusion Detection, Second Edition, Syngress.
[5] Angela Bonifati, Fabiano Cattaneo, Stefano Ceri, Alfonso Fuggetta, Stefano Paraboschi, Designing Data Marts for Data Warehouses in: ACM Transactions on Software Engineering and Methodology (TOSEM) Volume 10 , Issue 4 (October 2001), 452 – 483.
[6] Richard Bejtlich, "The Tao of Network Security Monitoring: Beyond Intrusion Detection", 2004 by Addison-Wesley.

[7] Cox, Kerry and Gerg, Christopher, 2004. Snort and IDS Tools, O'Reilly Media, Inc.
[8] Kimball, Ralph and Moss, Margy, 2002, The Data Warehouse Toolkit, Second Edition, John Wiley and Sons.
[9] Jacob Nikom , Real-time Sensor Data Warehouse Architecture Using MySQL, http://www.mysqluc.com/cs/mysqluc2005/view
[10] Snort Users Manual, Caswell, Brian and Hewlett, Jeremy, 2003, Snort.org.http://www.snort.org/docs/snort_htmanuals/htmanual_232
[11] Snort Database Plugin Documentation. Danyliw, Roman, 2002, Snort.org. August 2002. http://www.snort.org/docs/snortdb/snortdb.html
[12] Williams, Hugh E., and Lane, David, 2004, Web Database Applications with PHP and MySQL, O'Reilly Media
[13] E. Thomsen OLAP Solutions: Building Multidimensional Information Systems, Wiley, New York, NY (2000).
[14] Mizushima, M., Network Intrusion Detection System, Senior Project, 2005.
[15] Wasniowski R. Network Intrusion Detection System, RAW-RR-09-02, Report 2004.
[16] Wasniowski R., Multi-sensor agent-based intrusion detection system, Proceedings of the 2nd annual conference on Information security, Kennesaw, Georgia pp: 100 – 103, 2005.
[17] Michael Rash (Author), Angela D. Orebaugh, Graham Clark, Becky Pinkard, Jake Babbin Intrusion Prevention and Active Response: Deploying Network and Host IPS
[18] www.cipherdyne.org/**fwsnort**
[19] **http://www.securityfocus.com/infocus/1498**
[20] **http://www.techonline.com/community/related_content/21254**
[21] **http://www.mysqluc.com/cs/mysqluc2005/view/e_sess/6126**
[22] **http://www.honeypots.net/ids/links**
[23] **http://online.securityfocus.com/infocus/1498**
[24] **http://www.governmentsecurity.org/articles/HoneypotsDefinitionsandValueofHoneypots.php**