Chapter XXXVIII
An Overview of IDS Using Anomaly Detection

Lior Rokach
Ben-Gurion University of the Negev, Israel

Yuval Elovici
Ben-Gurion University of the Negev, Israel

ABSTRACT

Intrusion detection is the process of monitoring and analyzing the events occurring in a computer system in order to detect signs of security problems. The problem of intrusion detection can be solved using anomaly detection techniques. For instance, one is given a set of connection data belonging to different classes (normal activity, different attacks) and the aim is to construct a classifier that accurately classifies new unlabeled connections data. Clustering methods can be used to detect anomaly in data which might implies intrusion of a new type. This chapter gives a critical summary of anomaly detection research for intrusion detection. This chapter surveys a list of research projects that apply anomaly detection techniques to intrusion detection. Finally some directions for research are given.

INTRODUCTION

One of the most practical forms of cyber warfare is penetrating a mission-critical information system or any other critical infrastructure, and maliciously affecting its availability, confidentiality, or integrity. While the popularity of the Internet increases, more organizations are becoming vulnerable to a wide variety of cyber attacks. Thus, organizations employ various computer and network security solutions to make their information systems tolerant of such threats. One of the solutions is intrusion detection and prevention systems. Intrusion detection is the process of monitoring and analyzing the events occurring in a computer system and communication networks in order to detect signs of security breaches.

A complete *intrusion detection system* (IDS) might monitor network traffic, server and operating system events, and file system integrity, using both signature detection and *anomaly detection* at each level. Ma-

honey and Chan (2002) distinguish between a host based IDS, which monitors the state of the host and a network IDS, which monitors traffic to and from the host. These systems differ in the types of attacks they can detect. A network IDS can monitor multiple hosts on a local network. On the other hand, a host based system must be installed on the system it monitors. A host based system may, for example, detect user-to-root (U2R) attacks, where a certain user gains the privileges of another user (usually root). A network IDS detects probes (such as port scans), denial-of-service (DOS) attacks (such as server floods), and remote-to-local (R2L) attacks in which an attacker without user level access gains the ability to execute commands locally. Also, because a network IDS monitors input (and output) rather than state, it can detect failed attacks (e.g., probes).

There are two different approaches to intrusion detection: misuse detection and anomaly detection. Misuse detection is the ability to identify intrusions based on a known pattern for the malicious activity. These known patterns are referred to as signatures. These attack signatures encompass specific traffic or activity that is based on known intrusive activity. The reader is referred to the work of Axelsson (2000), for detailed taxonomy about IDSs.

The second approach, anomaly detection, is the attempt to identify malicious activity based on deviations from established normal activity patterns. Usually anomaly detection is performed by creating a profile for each user group. These profiles are used as a baseline to define normal user activity. If any monitored activity deviates too far from this baseline, then the activity generates an alarm.

Classic implementations of IDS are rule based (see Roesch, 1999). The system administrator is responsible to write a set of rules, for example, to reject any packet addressed to a nonexistent host, or to restrict services to a range of trusted addresses. However, keeping the rules updated by monitoring the traffic to determine normal behavior is challenging. Both types of intrusion detection systems can be benefit from using *data mining* techniques as will be shown later in the chapter.

BACKGROUND

Data mining is a term coined to describe the process of sifting through large and complex databases for identifying valid, novel, useful, and understandable patterns and relationships. Data mining involves the inferring of algorithms that explore the data, develop the model, and discover previously unknown patterns. The model is used for understanding phenomena from the data, analysis, and prediction. The accessibility and abundance of data today makes knowledge discovery and data mining a matter of considerable importance and necessity. Given the recent growth of the field, it is not surprising that a wide variety of methods is now available to researchers and practitioners.

Phung (2000) indicates that there are four shortfalls in classic IDS that data mining can be used to solve:

1. **Variants:** It is not uncommon for an exploit tool to be released and then have its code changed shortly thereafter by the hacker community. An example might be a Remote Procedure Call (RPC) buffer overflow exploit whose code has been modified slightly to evade an IDS using signatures. Since data mining is not based on predefined signatures the concern with variants in the code of an exploit are not as great.
2. **False positives:** A common complaint is the amount of false positives an IDS generates (i.e., alerting on non-attack events). A difficult problem that arises from this is how much can be filtered out without potentially missing an attack. With data mining it is easy to correlate data related to alarms with mined audit data, thereby considerably reducing the rate of false alarms (Manganaris, Christensen, Zerkle, & Hermiz, 2000) Moreover data mining can be used to tune the system and by that consistently reducing the number of false alarms.
3. **False negatives:** The dual problem of the false positive is the false negative in which an IDS does not generate an alarm when an intrusion is