Preface

Great changes are taking place is the area of information supply and demand due to the wide spread application of computers and the exponential increase of computer networks such as the Internet. The Internet has become a popular medium of commercial activities and this raised the stakes, both, for attackers and security personnel. Trillions of dollars of transactions occur daily at each major financial institution. For example, Visa processes 4,000 transactions per second, which means that if Visa's system goes down for one minute because of a distributed denial of service attack (DDoS), and assuming only $100 per transaction, over $24 million in transactions is lost in one minute.

Today the world of business computing is faced with the ever-increasing likelihood of unplanned downtime due to various attacks and security breaches. In this environment of uncertainty which is full of hackers and malicious threats, those companies around the globe which are the best at maintaining the continuity of their services (i.e., survive the system) and retaining their computing power, enjoy a significant competitive advantage.

Network downtime results in financial losses and more harms to the credibility of commercial enterprises especially ISPs. Minimizing or possibly eliminating the unplanned downtime of the system establishes the continuity of the computing services. Minimizing unexpected and unplanned downtime can be done by identifying, prioritizing and defending against misuse, attacks and vulnerabilities. The challenge is to reduce the likelihood of catastrophic incidents by: a) using appropriate machine and statistical learning techniques to assess the relative danger of individual threats and b) autonomously providing effective and appropriate response to the relevant threats.

Intrusion Detection System (IDS) is a rapidly growing field that deals with detecting and responding to malicious network traffic and computer misuse. Intrusion detection is the process of identifying and (possibly) responding to malicious activities targetd at computing and network resources. Any hardware or software automation that monitors, detects or responds to events occurring in a network or on a host computer is considered relevant to the intrusion detection approach. Different IDSs provide varying functionalities and benefits.

An attempt to break or misuse a system is called “intrusion”. An intrusion normally exploits a specific vulnerability and must be detected as quickly as possible. An intrusion detection system is a system for detecting such intrusions. Intrusion detection systems are notable components in network security infrastructure. They examine system or network activity to find possible intrusions or attacks and trigger security alerts for the malicious activities. They are generally categorized as signature-based and anomaly-based detection system. Other categories are network-based and host-based intrusion detection systems.

Network-based IDSs are placed at a strategic point or points within the network to examine passing network traffic for signs of intrusion, whereas, host-based IDSs are run on individual hosts or devices on the network and look at user and process activity for signs of malicious behavior.

In most networks, signature-based IDSs, which are very effective against known attacks, are deployed. Signature-based systems (also know as misuse detection systems) detect attacks based on known cases of misuse. Advantages of misuse detection are high confidence in detection, low false positive rate and an unambiguous detailed identification of attack. They are also more well understood and widely applied. Disadvantages are inability to detect unknown attacks and the need for expert knowledge to create signatures.

For defense against unknown attacks, an anomaly detection scheme has to be used, which creates a model of normal bhaviour of the system and detects deviation from this model. Techniques used in detecting anomalies include data mining, clustering, and statistical signal processing. The main advantage of anomaly based systems is the ability to detect unknown attacks. The disadvantages are high false positive rate and difficulty in identification of attack type. Moreover, since what is considered normal could be different in different environments, a distinct model of normalcy need to be learned individually.

A more recent class of intrusion detectors is the *specification-based* detectors, which try to reach a common ground between misuse-based and anomaly-based systems. They are mainly based on specifications derived from protocols and detect deviations from these specifications. Although they combine the benefits of anomaly detection and misuse detection, they suffer from the disadvantage that complete specifications are hard to create especially with most protocols being constantly extended.

Due to the exponential growth in size, distribution, and complexity of communication networks, current IDS technologies are not very effective against new attacks and have severe limitations as far as performance, scalability, and flexibility are concerned. Moreover, the improvements to the IDSs are often too slow and too little to keep up with the innovations by the attackers.

The main drawbacks of the current IDSs are: 1) the large number of false positives; 2) the inability to detect unknown attacks; and, 3) an inability to properly assess the relative danger of the misuse and provide an appropriate response. There is a general consensus that the primary focus of the intrusion detection technologies must be: a) to reduce the rate of false positives; b) to develop non-signature-based intrusion detection methods; and, c) work on prevention instead of detection.

There is a critical need to be able to deliver systems that can automatically detect intrusion patterns and performance bottlenecks, and dynamically defend themselves. The main goal of an intrusion detection system is to survive the system and retain its essential services. Survivability is often defined by resistance, recognition and recovery. Resistance deals with hardening a system to prevent a break-in or other malicious acts. The goal of recognition is to detect intrusive behavior from normal behavior. Recovery deals with ways of surviving malicious acts.

Any solution to the survivability problem must handle three basic criteria, including: dynamically changing network traffic patterns, occurrence of unpredictable events (security breaches), and non-finite base of network traffic environment. One way to develop survivability tools capable of blending seamlessly into current dynamic network environments is to design them as an agent-based distributed system. The key element of such a system is an intelligent agent, which is capable of analyzing a situation, making decisions and communicating with other agents and users. Multiagent systems together with the fuzzy systems can be used to establish a community of intelligent agents with features such as autonomy, self

During the past number of years, machine learning and data mining techniques have received considerable attention among the intrusion detection researchers to address the weaknesses of knowledge-base detection techniques. This has led to the application of various supervised and unsupervised techniques for the purpose of intrusion detection. This book provides a comprehensive review on current trends in intrusion detection systems and the corresponding technologies. We present a set of experiments which are carried out to analyze the performance of unsupervised and supervised machine learning techniques considering their main design choices.

During the last decade, anomaly detection has attracted the attention of many researchers to overcome the weakness of signature-based IDSs in detecting novel attacks. However, having a relatively high false alarm rate, anomaly detection has not been wildly used in real networks. This book presents data driven approaches to automating network behavior modeling. One of the approaches is the technique we developed to create an ARX model of network signals and using it for detecting network anomalies caused by intrusions. Network signals are nonstationary, highly volatile and hard to model using traditional methods. Our modeling technique using a combination of system identification theory and wavelet approximation is very effective at addressing this issue.

Alert correlation is an important technique for managing large volume of intrusion alerts that are raised by heterogeneous IDSs. The recent trend of research in this area is towards extracting attack strategies from raw intrusion alerts. Knowing the real security situation of a network and the strategies used by the attackers enables network administrators to launch appropriate response to stop attacks and prevent them from escalating. In this book we present alert management and correlation technique that can help to automatically extract attack strategies from a large volume of intrusion alerts without specific prior knowledge about these alerts.

The intrusion detection books on the market are relatively unfocussed. They tend to leave out details of a variety of key techniques and models. Additionally, many books lack much detail on different types of attacks, theoretical foundation of attack

detection approaches, implementation, data collection, evaluation, and intrusion response. In this book, our goal is to provide simple yet detailed and concise information on all these subjects. Additionally, we provide a detail overview of some of the commercially/publicly available intrusion detection and response systems. On the topic of intrusion detection system it is impossible to include everything there is to say on all subjects. However, we have tried to cover the most important and common ones.

This book is divided into 8 chapters and one appendix as follows:

Chapter 1 (Network Attacks): In this chapter we first discuss different attack taxonomies and then present the details of a large number of known vulnerabilities and strategies to lunch attacks.

Chapter 2 (Detection Approaches): Detection approach describes the attack analysis method used in an IDS. In this chapter different detection approaches that are currently available to detect intrusions and anomalous activities are given. Misuse, rule-based, model-based, anomaly and specification-based detection approaches are explained in detail.

Chapter 3 (Data Collection): Intrusion detection systems collect their data from various sources. Such sources include log files, network packets, system calls, or a running code itself. In this chapter, we provide detail information as to how this very important step in the life of different intrusion detection systems (i.e. host-based, network-based and application-based) can be accomplished.

Chapter 4 (Theoretical Foundation of Detection): Understanding the strengths and weaknesses of the machine learning and data mining approaches helps to choose the best approach to design and develop a detection system. Several approaches to the intrusion detection research area are introduced and analyzed in this chapter.

Chapter 5 (Architecture and Implementation): Intrusion detection systems can be classified based on their architecture and implementation. This classification usually refers to the locus of the data collection and analysis. This chapter introduces the centralized, distributed and agent based intrusion detection systems.

Chapter 6 (Alert Management and Correlation): Intrusion Detection Systems trigger too many alerts that usually contain false alerts. Decreasing false positives and improving the knowledge about attacks provides a more global view of what is happening in a network. Alert management and correlation addresses the issue of managing large number of alerts by providing a condensed, yet more useful view of the network from the intrusion standpoint. The correlation function can relate different alerts to build a big picture of the attack. The correlated alerts can also be used for cooperative intrusion detection and tracing an attack to its source.This chapter introduces different approaches to cluster, merge and correlate alerts.

Chapter 7 (Evaluation Criteria): This chapter provides a number of approaches that can be used to evaluate the potential intrusion detection systems for accuracy, performance, completeness, timely response, cost and intrusion tolerance & attack resistance.

Chapter 8 (Intrusion Response): The objective is to enable automated reasoning and decision-making to aid in human-mediated or automatic response. The cost-benefit analysis of response is the key for effective response. This chapter expands on this and other approaches for providing effective and sensible responses.

Appendix A (Examples of Commercial and Open Source IDSs): A brief introduction to some of the current commercial and open source IDSs are given in Appendix A.

Audience: This book provides students and security professionals with the necessary technical background to understand, develop and apply intrusion detection systems. Some background of machine learning and data mining is helpful to understand the approaches presented in this book. We also assume that the readers of this book have a good command of data communication and networking. After reading this book, you should have a solid understanding of technical basics of intrusion detection and response systems. In addition, you should know how to develop reliable and effective detection systems and be able to install and manage commercially available IDSs.

Fredericton, Canada,
July 2009

Ali A. Ghorbani
Wei Lu
Mahbod Tavallaee

http://www.springer.com/978-0-387-88770-8

Network Intrusion Detection and Prevention
Concepts and Techniques
Ghorbani, A.A.; Lu, W.; Tavallaee, M.
2010, XVIII, 216 p. 20 illus., Hardcover
ISBN: 978-0-387-88770-8