Abstract

Anomaly-based intrusion detection is to discern malicious and legitimate patterns of behavior in the variables characterizing the normality of information systems. As the system normality is constructed only from an observed sample of normally occurring patterns, despite systems are dynamic in nature (or user-driven) and becoming increasingly complex, anomaly detectors often suffer from excessive false alerts. This dissertation presents our work on the development of effective and efficient models, methods and techniques for anomaly-based intrusion detection with the main concern of adaptability. Our work generally involves three parts, which can be summarized as follows:

The first part of our work is motivated by the observation that the fundamental understanding of the computing environments is an initial but essential step in the process of developing an effective anomaly detection model. Based on the similarity between anomaly detection and induction reference problems, we present a statistical framework for analyzing the general behavior of anomaly detectors from the perspective of their observable subjects. The objective is to lay a theoretical foundation for the modeling and developing of our specific anomaly detectors in the next stage of our work. To enrich the framework, we examine some challenging issues currently exist and present potential solutions, including host-based and network-based normality characterization, evaluation of anomaly detectors, etc.; The framework also involves some case studies and comparative analysis on several typical anomaly detectors, for the sake of presenting a formal way for the understanding and development of anomaly detectors' operational characteristics, and therefore bring them to broader applications.

Taking the constructed framework as starting point, and with the objective to capture the normality drifts of computer systems behavior driven by users or system itself, we develop three versions of SVM-based anomaly detectors, which employ three modified Support Vector Machines as the kernel detection scheme. Our modification aims to break the traditional assumption that anomaly detectors are always fed with training data that are readily available with desired quality in batch, and thus enable them to be trained online periodically for the sake of adapting to the new computing environments without triggering excessive false alerts. To validate those anomaly detector's performance, we implement the experiments by reforming 1998 DARPA BSM data set collected at MIT's Lincoln Labs, and conduct the comparative studies with the original algorithms. The experimental results verify that our new designed anomaly detectors outperform the original ones with fewer support vectors (SVs) and less training time without sacrificing detection accuracy.

Based on the observations and conclusions of the first part of work, we present another framework for the correlation of several observation-specific anomaly detectors. Our hope is that a collection of simple surrogates based on specific operating environments can cooperate well and evolve into generic models with broader anomaly detection coverage and less false alerts. As the specific implementation of the framework, we develop an integrated anomaly detection model named Autonomic Detection Coordinator (ADC) to defend against host-based intrusive anomalies. The cooperation between four host-based anomaly detectors is formulated as a multi-agent partially observable markov decision process. A policy-gradient reinforcement learning algorithm is then employed to search in an optimal cooperation manner, with a set of parameters controlling individual anomaly detector's behavior. The generic behavior of the coordinator can be adjusted easily by setting a reward signal to adapt to changing system environments. A host-based experimental scenario is developed for implementation, and the experimental results show its satisfactory performance. The model is also extended as the basic framework for the modeling and analysis of multi-stage coordinated attacks in computer networks. The definitions and properties derived from the models (both defender-centric and attacker-centric) present us a formal way for the development of countermeasures to thwart or mitigate such attacks. Taking into account the specific concerns of attackers and defenders, two algorithms called Attackers Nondeterministic Trail Searching algorithm (ANTS) and Attacker's Pivots Discovery by Backward Searching algorithm (APD-BS) are developed respectively. The former one aims to search for the most efficient concurrent actions for attackers, and the latter one intends to discover the attacker's significant observations for defenders.