Survey of Current Network Intrusion Detection Techniques

Sailesh Kumar, sailesh@arl.wustl.edu

Abstract:

The importance of network security has grown tremendously and a number of devices have been introduced to improve the security of a network. Network intrusion detection systems (NIDS) are among the most widely deployed such system. Popular NIDS use a collection of signatures of known security threats and viruses, which are used to scan each packet's payload. Signature based designs have low false positive rates, and they are effective and accurate in combating against the known security threats. However, they remain completely ineffective against those attacks that are yet unknown; these can be combated only after they are detected manually and a signature is created for them.

Since new threats are potentially more lethal, a number of pro-active designs have been proposed, which can detect new security events such as propagation of a new and unknown virus or worm. Such systems accomplish this by creating a profile of normal Internet traffic, and then using this profile to continuously monitor the network activity for suspicious activity. As the system senses an anomaly, or a dramatic change in traffic characteristics, it takes certain actions such as raising an alarm or discarding certain traffic. In this Survey paper, we will evaluate a number of current NIDS systems and the algorithms they employ to detect and combat security threats, both from technical and economical perspective.

Keywords:

NIDS, Anomaly Detection, Network Security, Security Signature, Pattern Matching

Table of Contents

---

1. Introduction

Network security has recently received an enormous attention due to the mounting security concerns in today's networks. A wide variety of algorithms have been proposed which can detect and combat with these security threats. Among all these proposals, signature based Network Intrusion Detection Systems (NIDS) have been a commercial success and have seen a widespread adoption. While, these systems already generate several hundreds of million dollars in revenue, it is projected to rise to more than 2 billion dollars by 2010.

A NIDS aims at detecting possible intrusions such as a malicious activity, computer attack and/or computer misuse, spread of a virus, etc, and alerting the proper individuals upon detection. A NIDS monitors and analyzes the data packets that travel over a network looking for such suspicious activities. A large NIDS server can be set up on the links of a backbone network, to monitor all traffic; or smaller systems can be set up to monitor traffic directed to a particular server, switch, gateway, or router. Another class of NIDS can be setup at a centralized server, which will scan the system files, looking for unauthorized activity and to maintain data integrity.

There are two primary approaches to NIDS implementation: signature based, and anomaly detection based. The first approach has become a commercial success. A signature based NIDS maintains a collection of signatures, each of which characterizes the profile of a known security threat (e.g. a virus, or a DoS attack). These signatures are used to parse the data streams of various flows traversing through the network link; when a flow matches a signature, appropriate action is taken (e.g. block the flow or rate limit it). Traditionally, security signatures have been specified as a string signature, port signature and header condition signature.

String signatures are a string of ASCII symbols that characterizes a known attack. For example, such a string signature in UNIX can be "cat "+ +" > /.rhosts" , which if executed, can cause the system to become extremely vulnerable to network attack. Simpel strings may lead to high false positives, therefore it is important to refine the string signature; for this purpose one may use a compound string signature. Such a compound string signature to detect a common Web server attack can be "cgi-bin" AND "aglimpse" AND "IFS".

Port signatures commonly probes for the connection setup attempts to well known, and frequently attacked ports. Obvious examples include telnet (TCP port 23), FTP (TCP port 21/20), SUNRPC (TCP/UDP port 111), and IMAP (TCP port 143). If these ports aren't being used by the site at a point in time, then the incoming packets directed to these ports are considered suspicious.

Header signatures are designed to watch for dangerous or illegitimate combinations in packet headers fields. The most famous example is Winnuke, in which a packet's port field is NetBIOS port and one of the Urgent pointer, or Out Of Band pointer is set. In earlier version of Windows, this resulted in the "blue screen of death". Another well known such header signature is a TCP packet header in which both the SYN and FIN flags are set. This signifies that the requestor is attempting to start and stop a connection simultaneously.

Some well known commercial NIDS include AXENT (www.axent.com), Cisco (www.cisco.com), CyberSafe (www.cybersafe.com), ISS (www.iss.net), and Shadow (www.nswc.navy.mil/ISSEC/CID), while the popular open source NIDS includes Snort, and Bro.

While signature based NIDS has been widely deployed, anomaly based NIDS have not gained popularity yet, and they

have remained a topic of an ongoing interest among the research community. The critical advantage of such NIDS over signature based NIDS is its promise to detect and contain security violations before they propagate and cause any damage. Signature based systems are reactive, in that they combat against known attacks, that have already affected and damaged a number of systems before being identified; anomaly based systems are pro-active and autonomous and can ensure security without any manual interference.

Anomaly based NIDS monitors network traffic and compares it against an established baseline of normal traffic profile. The baseline characterizes what is "normal" for the network - such as the normal bandwidth usage, the common protocols used, correct combinations of ports numbers and devices - and alerts the administrator or user anomalous traffic is detected which is significantly different from the baseline. It is highly subjective to decide what can be considered normal and what an anomaly, but a widely accepted rule of thumb is that, any incident which occurs on a frequency greater than two standard deviations from the statistical norm should be considered suspicious. An example of such behavior would be if a normal user logs on and off of a machine 20 times a day instead of the normal course of 1 or 2 times. Anther example is, when a user computer is used at 2:00 AM at a time when no one outside of the business hours have access; this should also raise some suspicions. At another level, a NIDS can investigate the user patterns, such as profiling the programs that are often executed, etc. If a user in the administrative department suddenly starts to execute programs from the engineering division, or begins to compile a code, then the system can promptly alert the administrators.

Clearly, such anomaly based intrusion detection may lead to a high rate of false detection, which we call false positives. It is generally considered difficult to keep low false positives in any system that sets aggressive policies to detect anomalies. For example, it may be difficult to distinguish flash crowd from a Distributed Denial of Service attack (DDoS), thus a system may raise false alarm during a flash crowd event assuming that it is a DDoS attack. Similarly network reconfigurations and transient failures may abruptly change the traffic profile falsely raising the alarm. The second challenge concerns with the assumption made by these systems that attacks are always anomalous, which may not necessarily be true. An intelligent attacker may develop intrusion techniques which will cause minimal disruption in the underlying traffic, thus may go undetected.

The final challenge in designing these systems concerns with the availability of dataset that is representative of normal traffic. To be realistic, the assumption that there exists attack-free data for training a detector outside of simulated data is not a realistic assumption. Typical network traffic contains a large number of scans, denial-of-service attacks and backscatter, and worm activity. If not careful, this activity will become part of the normal state for an anomaly detector.

In this survey paper, we describe the in-depth design details of both signature- and anomaly based NIDS. Specifically, for signature based systems we cover the following topics:

1. An overview of the system - high level architecture and principles of intrusion detection.
2. Current well known systems and the algorithms and architectures employed by them.
    a. Signatures that are used.
    b. Hardware/system architecture.
    c. Performance and limitations.
3. Effectiveness in improving the network security.
4. The disadvantages of such systems.

For the anomaly based systems, we cover the following topics:

1. Overview of anomaly detection systems, and high level architecture.
2. Well known algorithms that are used to detect anomalous behavior within well behaved traffic.
    a. Unsupervised clustering algorithms.
    b. Entropy based methods.
    c. Data mining techniques.
3. How anomaly detection can be used to combat security threats, and the discussion of their possibly important role in the near future.
4. The disadvantages and challenges in implementing such systems and a brief description of the techniques that are commonly used by attackers to evade them.

Back to Table of Contents
