Digital Investigation 20 (2017) S56–S65

Contents lists available at ScienceDirect

Digital Investigation

journal homepage: www.elsevier.com/locate/diin

DFRWS 2017 Europe — Proceedings of the Fourth Annual DFRWS Europe

Evidence gathering for network security and forensics

Dinil Mon Divakaran*, Kar Wai Fok, Ido Nevat, Vrizlynn L.L. Thing

*Cyber Security Cluster, A*STAR Institute for Infocomm Research (I²R), 1 Fusionopolis Way, #21-01 Connexis (South Tower), 138632, Singapore*

ARTICLE INFO

Article history:
Received 31 January 2017
Accepted 31 January 2017

Keywords:
Forensics
Security
Network
Traffic
Regression

ABSTRACT

Any machine exposed to the Internet today is at the risk of being attacked and compromised. Detecting attack attempts, be they successful or not, is important for securing networks (servers, end-hosts and other assets) as well as for forensic analysis. In this context, we focus on the problem of evidence gathering by detecting fundamental patterns in network traffic related to suspicious activities. Detecting fundamental anomalous patterns is necessary for a solution to be able to detect as many types of attacks and malicious activities as possible. Our evidence gathering framework correlates multiple patterns detected, thereby increasing the confidence of detection, and resulting in increase in accuracy and decrease in false positives. We demonstrate the effectiveness of our framework by evaluating on a dataset consisting of normal traffic as well as traffic from a number of malwares.

Introduction

End-hosts and networks become prone to attacks and compromises once they get connected to the Internet. The types and sophistication of attacks have only been increasing with time, one of the motivations being the monetary profits that cyber-criminals obtain from such malicious and cyber–criminal activities. To quantify, in its annual security report of 2016, Cisco's estimation of the average annual income for ransomware per campaign is $34 million (Cisco Annual Security Report, 2016).

Detection of anomalies related to attacks is imperative for securing a network as well as for forensic analysis. Timely detection can thwart damages due to the attack; forensic analysis of the anomalies aid not only in investigations but also in profiling of attacks. There exist a plethora of research works on detecting anomalies and attacks in networks (Carl et al. (2006), Feily et al. (2009), Chandola et al. (2009), Vincent Zhou et al. (2010) and Bhuyan et al. (2014) are some relevant surveys). Often, different solutions are developed for different attacks, leading to complex management and higher costs for the users. Besides, individual anomalies are usually detected independently; and without having a mechanism to check if some of the detected anomalies are correlated, solutions could easily miss the forest for the trees. Worse, the detected anomalies might be assigned a low risk, and therefore ignored by an analyst.

Differing from the previous works, our aim is to develop a framework that can gather as many evidences as possible for different types of attacks and malicious activities. In this work, we focus on the problem of evidence gathering by analyzing network traffic. The evidences are the fundamental patterns related to suspicious activities. Detecting such patterns allows for detection of anomalies that are common to a number of attacks. We take a pragmatic approach in detecting anomalous patterns; we apply regression models to analyze network traffic data and detect patterns related to suspicious activities. The anomalous patterns, which form evidences, are correlated and analyzed, to increase the performance of detection of traffic related to attacks and malicious activities, where performance is expressed in terms of detection accuracy and false positive rate.

The advantages of our proposed system are:

1. No learning of normal behavior using benign data: Given that Internet traffic is always evolving and changing, a solution that models static data will not be adaptive. Besides, it is extremely difficult to obtain a real-life up-to-date traffic dataset of meaningful size that consists of only benign flows and captures the dynamic characteristics of benign traffic. Our approach, on the other hand, uses information of only attack traffic. Characteristics of attack traffic are very specific to the kind of attack, and also differs significantly from that of normal traffic.
2. Detection of multiple anomalous patterns: By detecting anomalous patterns in network traffic, our solution essentially gathers evidences related to the attack. For a single attack, there might be multiple patterns that are detected, some due to different

* Corresponding author.
E-mail addresses: divakarand@i2r.a-star.edu.sg (D.M. Divakaran), fokkw@i2r.a-star.edu.sg (K.W. Fok), ido.nevat@tum-create.edu.sg (I. Nevat), vriz@i2r.a-star.edu.sg (V.L.L. Thing).

http://dx.doi.org/10.1016/j.diin.2017.02.001

features used for analysis, while some due to the different stages of an attack (such as, reconnaissance, penetration and exploit). More the number of patterns, higher the confidence of detection.
3. Detection of multiple attacks: As our solution searches for different patterns related to attacks, it is not tailored for detection of any particular attack; in other words, it is designed to detect different kinds of attacks. As we demonstrate in Section "Performance evaluations", our solution is able to detect traffic generated by a number of different malwares.

We evaluate our solution on a dataset compiled from different sources, and consisting of both normal traffic as well as traffic generated by malwares. The evidence-gathering framework we developed is able to detect malware traffic sessions with high accuracy, while maintaining low false positive rate. Different from past works, our evaluation also provides insights into the number of evidences detected for malicious traffic sessions, thereby aiding operators to decide on the configurations.

We discuss related works in the next section. In Section "Framework for evidence gathering", we present the framework for evidence gathering. The techniques based on regression modeling and analysis that we use for detection of anomalous patterns are presented in Section "Regression analysis for detection of anomalous patterns". The eventual decision making criteria, on whether a (set of) traffic flows anomalous or not, is based on the evidences gathered, and this is presented in Section "Decision making based on evidences gathered". We present the experiments performed and discuss the results obtained in Section "Performance evaluations".

Related works

Tremendous amounts of efforts have been invested by the research community to develop solutions for detection of network attacks and anomalies. Even the literature on one specific set of attacks, for instance DoS attacks (Carl et al., 2006; Peng et al., 2007), is vast. This section provides a brief discussion on some of the important related works, without attempting to be exhaustive.

Although the terms anomalies and attacks are highly related, there is also a difference. An anomaly is anything that deviates from what is defined and observed as normal behavior. The cause of such deviations might be network faults (e.g., link failure), abrupt changes in network—for example route changes, or even changes in CDN (Content Delivery Network) caches (Fiadino et al., 2014), attacks, etc. While attacks are all activities that (attempt to) breach the security of any given system, not all anomalies are attacks. Examples of such anomalies (that are not attacks) range from sudden peak in traffic at a web server to receiving large number of scanning packets on a network. Yet, a port scan on a machine might be related to an impending penetration attempt or vulnerability exploitation at a specific port. Therefore, while individual anomalies may well turn out to be evidences of attacks and compromises, analyzing them independently might lead to unacceptable false-positive rate, thereby becoming counterproductive.

To solve the challenging problem of anomaly detection in networks, researchers have applied knowledge from different (overlapping) spheres such as expert system (Ilgun et al., 1995), information theory (Gu et al., 2005; Nychis et al., 2008), machine learning and data mining (Lee and Stolfo, 1998; Portnoy et al., 2001; Lakhina et al., 2005), signal processing (Barford et al., 2002; Thottan and Ji, 2003), statistical analysis (Kruegel et al., 2003; Simmross-Wattenberg et al., 2011; Thatte et al., 2011) and pattern recognition (Fontugne and Fukuda, 2011). We refer to a recent survey for discussions on the applications of some of these approaches (Bhuyan et al., 2014).

A rule-based system to detect penetrations by modeling them as state transitions was proposed in Ilgun et al. (1995). But, this was developed for the Unix system where actions or state changes occur on the execution of commands by the attacker. Besides, rule-based systems do not adapt to the fast evolving nature of network traffic.

Entropy has been used to evaluate features for network anomaly detection, in particular to understand the correlation of different features (both header features and behavioral features), emphasizing the need to select features with care (Nychis et al., 2008). One of the well-known works which builds on the principle of maximum entropy to detect anomalies in network is developed in Gu et al. (2005). Features are added one by one iteratively, such that in each step, the right weight for each feature is estimated using KL (Kullback–Leibler) divergence. Learning the parameters and thereby building a model from a given data, KL divergence is again used to detect anomalies. Clustering based anomaly detection solutions that used unlabeled data have also been proposed in the past (Portnoy et al., 2001; Leung and Leckie, 2005; Jiang et al., 2006). For example, Portnoy et al. (2001) develops a variant of single-linkage clustering to build normal clusters in unlabeled data, and subsequently use the clusters to detect anomalies.

A number of statistical techniques have also been explored in the past. In Simmross-Wattenberg et al. (2011) traffic was modeled using α-stable distributions, and anomalies such as flash and flash crowds are detected by comparing trained traffic windows with the test windows using the Generalized Likelihood Ratio Test (GLRT). This solution depends on labeled data; in addition, it is also computationally expensive.

More recent works explore traffic at multiple time resolutions and also deploy an ensemble of techniques. One such work is the unsupervised anomaly detection system developed in Casas et al. (2012); time-series are build at different time scales, and outlier traffic flows are ranked after they are identified using a combination of cluster techniques (namely, Sub-Space Clustering, Density-based Clustering and Evidence Accumulation Clustering). While the solution is shown to detect anomalies and attacks, the ability to detect more sophisticated attacks beyond DoS and DDoS attacks is not known.

Different from the above works, our focus is explicitly on gathering evidences by modeling and analyzing network traffic. By gathering multiple evidences, the solution is able to reduce false positives to acceptable levels. The framework we develop also does not require learning or building models based on normal traffic. Besides, our solution is more general in design (and not tailored for any specific attack), as the patterns being detected are fundamental to different malicious activities.

Framework for evidence gathering

Fig. 1 illustrates the framework that we develop for gathering evidences of and related to attacks in network traffic. Our system takes the following three-stage approach to detect attack sequences:

1. Model sessions of traffic flows, and detect anomalous patterns.
2. Detect network and port scans, as well as illegitimate TCP state sequences.
3. Gather and correlate anomalous patterns, and make final decision on traffic (whether it is anomalous or not).

We explain each of the above stages below.

Stage 1: Modeling and analyzing sessions to detect anomalous patterns

To start with, we define flows and sessions. A flow is a set of packets, localized in time, with the same five tuple of source and