Available online at www.sciencedirect.com

ScienceDirect

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

Computers & Security

An anomaly analysis framework for database systems

Sokratis Vavilis [a,*], Alexandru Egner [a], Milan Petković [a,b], Nicola Zannone [a]

[a] Eindhoven University of Technology, Den Dolech 2, Eindhoven 5612AZ, Netherlands
[b] Philips Research Eindhoven, High Tech Campus 34, Eindhoven 5656AE, Netherlands

ARTICLE INFO

Article history:
Received 29 October 2014
Received in revised form
28 April 2015
Accepted 5 June 2015
Available online 2 July 2015

Keywords:
Anomaly detection
Data leakage
Risk assessment
Database attack classification
Alert visualization

ABSTRACT

Anomaly detection systems are usually employed to monitor database activities in order to detect security incidents. These systems raise an alert when anomalous activities are detected. The raised alerts have to be analyzed to timely respond to the security incidents. Their analysis, however, is time-consuming and costly. This problem increases with the large number of alerts often raised by anomaly detection systems. To timely and effectively handle security incidents, alerts should be accompanied by information which allows the understanding of incidents and their context (e.g., root causes, attack type) and their prioritization (e.g., criticality level). Unfortunately, the current state of affairs regarding the information about alerts provided by existing anomaly detection systems is not very satisfactory. This work presents an anomaly analysis framework that facilitates the analysis of alerts raised by an anomaly detection system monitoring a database system. The framework provides an approach to assess the criticality of alerts with respect to the disclosure of sensitive information and a feature-based classification of alerts according to their associated type of attack. The framework has been deployed as a web-based alert audit tool that provides alert classification and risk-based ranking capabilities, significantly easing the analysis of alerts. We validate the classification and ranking approaches using synthetic data generated through an existing healthcare management system.

1. Introduction

Database management systems (DBMS) are used to collect, store, disseminate, and analyze sensitive data such as customer records and confidential business information. This makes DBMS an attractive target for attackers. A study conducted by Verizon in 2014 shows that DBMS are one of the most compromised assets of organizations (Verizon, 2014). Therefore, they are an essential asset that needs to be protected. However, it is very difficult, if not impossible, to prevent database attacks completely. Thus, prevention solutions are often coupled with detection and incident response management solutions. These solutions, however, should satisfy stringent constraints on how security incidents are handled. For instance, the new EU data protection regulation obliges organizations to take actions and notify data protection authorities within 24 h after the occurrence of a data breach is detected (Information Age, 2012). Thus,

* Corresponding author.
E-mail addresses: s.vavilis@tue.nl (S. Vavilis), a.i.egner@tue.nl (A. Egner), milan.petkovic@philips.com (M. Petković), n.zannone@tue.nl (N. Zannone).
http://dx.doi.org/10.1016/j.cose.2015.06.004
0167-4048/

organizations need solutions that allow *timely* detection and response to security incidents.

The detection of security incidents in DBMS is often addressed by employing anomaly detection systems (Costante et al., 2014; Kamra et al., 2008; Mathew et al., 2010; Mohammadian and Hatzinakos, 2013; Shebaro et al., 2013). Anomaly detection systems monitor user database activities (defined in terms of database queries) and raise alerts when suspicious activities are detected. Various techniques have been proposed to detect anomalies (see Chandola et al. (2009) for a survey). In broad terms, anomaly detection systems can be classified into rule-based and behavior-based. Rule-based approaches (Hart et al., 2011; Takebayashi et al., 2010) use predefined rules to specify activity patterns that correspond to anomalous behavior. These approaches, however, are only able to detect known attacks, leaving unknown attacks undetected. This issue is addressed by behavior-based approaches (Costante et al., 2014; Koch, 2011), which automatically learn user profiles by observing "normal" activities and marking every deviation from such profiles as a potential threat.

The alerts raised by an anomaly detection system should be investigated before deciding the actions to be taken in order to respond to a security incident (Tndel et al., 2014). There are two main challenges to be faced when analyzing alerts:

- *Gathering of alert information*: Deciding how to respond to a security incident requires a deep understanding of the alerts. Such an understanding is only possible if relevant information about alerts, such as *root causes* and *attack type*, are known. However, gathering such information manually is costly and time-consuming.
- *Large number of alerts*: The number of alerts raised by the anomaly detection system can be huge. For example, a large number of alerts can be raised in hospitals due to the usage of the break-the-glass protocol (Banescu and Zannone, 2011). In addition, existing anomaly detection systems and, in particular, behavior-based anomaly detection systems are often characterized by a high rate of false positives (Hadiosmanovi et al., 2012; Santos et al., 2014), i.e. alerts that are not actual security incidents. Identifying the most critical incidents among a large number of alerts requires substantial efforts.

To facilitate the analysis of alerts, it is thus desirable to have means for (*i*) gathering the information necessary for the analysis and (*ii*) prioritizing alerts with respect to their criticality.

Unfortunately, most existing anomaly detection systems provide no or very limited information about alerts (Costante et al., 2013), thus failing to meet desideratum (*i*). Few proposals (Costante et al., 2014; Mathew et al., 2010) give intuition on how to enrich alerts with relevant information, but do not provide a systematic approach for gathering such information. Some approaches in the field of anomaly-based intrusion detection (Bolzoni et al., 2009; Goseva-Popstojanova et al., 2014; Gowrison et al., 2013) provide attack classifications of alerts. However, these approaches are mainly tailored for the analysis of network traffic and are not suitable for the analysis of database activities.

To identify the most severe security incidents (desideratum (*ii*)), few approaches provide support for the *quantification* and *ranking* of alerts based on their criticality. In particular, some approaches (Costante et al., 2014; Viswanathan et al., 2012) quantify alerts with respect to their anomaly level, allowing security officers to focus on the most abnormal cases. Other approaches (Harel et al., 2012; Vavilis et al., 2014) aim to provide an estimation of the damage caused by data leakages on the basis of the sensitivity and amount of leaked information, allowing security administrators to focus on the most severe data breaches. These approaches, however, only provide a "partial view" on the criticality of alerts, which may lead to a ranking that does not reflect their actual criticality level. Indeed, the criticality of security incidents is typically measured in terms of *risk*, where risk is defined as a function of the likelihood (probability) of an undesired event occurring and the impact (severity) that such an event has on an organization (ISO/IEC, 2009).

In this work, we propose a unified framework to facilitate the analysis and evaluation of alerts raised by an anomaly detection system monitoring user database activities. The framework aims to enrich alerts with information necessary for their analysis. First, we propose a risk-based data leakage quantification approach that allows an estimation of the criticality of alerts with respect to the disclosure of sensitive data. To provide a complete view on the criticality of alerts our risk metric encompasses both the probability of an alert to be a true alert and its severity with respect to amount and sensitivity of the data leaked. For estimating the probability that an alert corresponds to a true security incident, we adopt and extend the white-box behavioral-based anomaly detection system presented in Costante et al. (2014). This system provides an anomaly score together with the alerts raised, which indicates whether they are a real threat or not. In our work, we revise how the anomaly score is computed in order to express such a score as a probability. To calculate the severity of data leakages, we employ the data leakage quantification approach presented in Vavilis et al. (2014). We refer to Section 2 for an overview of these two approaches. Second, we propose a novel method for the classification of alerts with respect to database attacks. We study the most frequent database attacks (Imperva, 2014) and elicit the features characterizing these attacks. We show how these features can be used to specify rules representing attack patterns. Based on these rules, we present an approach to classify alerts with respect to database attacks. To enable the timely analysis and evaluation of alerts, we present a database anomaly audit tool for the visualization of alerts. In particular, the tool shows in a user-friendly manner the relevant information about alerts, including their criticality level (with respect to data leakage) and attack type. Thus, the tool facilitates alert analysis by making it possible to focus on the most critical data leakages or on specific types of database attacks.

We validate the risk-based data leakage quantification approach and the feature-based attack classification method using a case study in the healthcare domain. Healthcare is indeed an interesting domain to investigate as a large amount of sensitive data, such as patient healthcare records, has to be protected. For the experiments, we have generated, together with our industry partner, synthetic datasets of database