Information Fusion 10 (2009) 285–299

Contents lists available at ScienceDirect

Information Fusion

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

A logic-based model to support alert correlation in intrusion detection

Benjamin Morin [a,*], Ludovic Mé [a], Hervé Debar [b], Mireille Ducassé [c]

[a] *Supelec, Avenue de la Boulaie, CS-47601, F-35576 Cesson-Sevigné, France*
[b] *Orange Labs, 42 rue des Coutures, BP-6243, F-14066 Caen, France*
[c] *IRISA, Campus de Beaulieu, F-35042 Rennes, France*

ARTICLE INFO

Article history:
Received 21 December 2006
Received in revised form 17 March 2008
Accepted 23 January 2009
Available online 7 February 2009

Keywords:
Intrusion detection
Alert correlation
Data model

ABSTRACT

Managing and supervising security in large networks has become a challenging task, as new threats and flaws are being discovered on a daily basis. This requires an in depth and up-to-date knowledge of the context in which security-related events occur. Several tools have been proposed to support security operators in this task, each of which focuses on some specific aspects of the monitoring. Many alarm fusion and correlation approaches have also been investigated. However, most of these approaches suffer from two major drawbacks. First, they only take advantage of the information found in alerts, which is not sufficient to achieve the goals of alert correlation, that is to say to reduce the overall amount of alerts, while enhancing their semantics. Second, these techniques have been designed on an ad hoc basis and lack a shared data model that would allow them to reason about events in a cooperative way. In this paper, we propose a federative data model for security systems to query and assert knowledge about security incidents and the context in which they occur. This model constitutes a consistent and formal ground to represent information that is required to reason about complementary evidences, in order to confirm or invalidate alerts raised by intrusion detection systems.

1. Introduction

Managing and supervising security in large networks has become a challenging task, as new threats and flaws are being discovered on a daily basis. This requires an in depth and up-to-date knowledge of the context in which security-related events occur. Several systems have been proposed to support security operators in this task. Some examples of these systems include:

- intrusion detection systems (IDS), which monitor the activity of the information system for the occurrence of malicious activities,
- firewalls, which filter inbound and outbound network traffic,
- vulnerability assessment scanners, which discover and report potential risks in computer systems,
- active and passive network mapping systems, which provide a picture of the network nodes and their interconnections (so called topology), as well as the software products running on them (so called cartography),
- honey-pots, which report current trends and threats in the wild and incident databases which inventory attack characteristics.

The scope of each of these systems is limited, both in terms of detection capabilities and in the part of the network they monitor. Therefore, several sensors need to be dispatched throughout the network in order to provide security operators with a comprehensive view of the events that occur. Since each system is likely to produce a large amount of observations, many of which are incomplete, irrelevant or not reliable, security operators are rapidly overwhelmed with events, the analysis of which is complex and time consuming. Thus, it is necessary to assist security operators in the diagnosis of the security incidents in order for them to focus on high priority incidents and take appropriate counter-measures.

Therefore, several event correlation and reasoning approaches have been proposed in the literature to fuse information available in alerts triggered by security mechanisms. However, information available in alerts is generally not sufficient; information about the monitored information system, the characteristics of attacks, and the configuration of security devices deployed throughout the network are also necessary. We argue that having a common data model to describe the relevant security-related information is a prerequisite for the security systems to share a common understanding of the situation at stake, and cooperate.

For this purpose, we propose a data model based on first order logic for security systems to query and assert information about security incidents and the context in which they occur. This model constitutes a consistent and formal ground to represent knowledge that is required to reason about complementary intrusion evidence.

[*] Corresponding author. Tel.: +33 2 99 84 45 78.
E-mail addresses: benjamin.morin@supelec.fr (B. Morin), ludovic.me@supelec.fr (L. Mé), herve.debar@orange-ft.com (H. Debar), mireille.ducasse@irisa.fr (M. Ducassé).

1566-2535/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.inffus.2009.01.005

This paper follows prior work on modeling knowledge in the intrusion detection field. In [1], we proposed a relational data model called M2D2 (from the initials of the authors' name), whose objective was to federate the information that is required to reason about alerts in intrusion detection. M2D2 was a first attempt to put together the concepts of alerts, events, vulnerabilities, sensors, network hosts, software products and their relationships.

Our contribution in this paper is twofold. Firstly, we have completely reformulated M2D2 in the first-order logic formalism. This new formalism allows one to translate almost straightforwardly the modelled concepts and relationships in an operational system, thanks to the existing prolog and datalog systems. This revised formalism also allows us to take advantage of logic as a uniform language to represent knowledge databases. Facts, rules and queries can be written in a single language. Moreover, the formalism supports definition of relations via recursive rules, something which is not allowed in traditional databases. Secondly, the new model includes new concepts that were missing in the original M2D2 model. In addition, the modeling of some remaining concepts has been refined. These changes include the description of attack instances and classes, a finer integration of security system capabilities, and the taking into account of routing in networks.

This article is structured as follows. First, we discuss the background and motivation of our work. The next four sections focus on each family of concepts our model is made of: the context (i.e. the characteristics of the monitored information system), the attacks and vulnerabilities, the security devices and the events and alerts that occur in the system. In Section 7, we show how the model can be used to reason about alerts by means of an attack scenario. Then, we briefly describe how a prototype implementation of the logical framework fits within an alert management platform. Before concluding and discussing future work, we present some related work on the subject.

2. Background and motivations

2.1. Intrusion detection

Intrusion detection is a field of computer security whose goal is to monitor the activity of an information system for the occurrence of malicious activities, i.e., actions intended to violate the security policy governing confidentiality, integrity and availability of services and data.

Intrusion detection has been a very active research area for the past 20 years, and several complementary solutions have been proposed to detect attacks of all forms and origins against hosts and networks. Despite these efforts, intrusion detection systems (IDS) still suffer from several drawbacks. Firstly, IDS trigger too many alerts, a large proportion of which turn out to be false or irrelevant alerts [2]. Security operators are consequently overwhelmed with alerts, the analysis of which is time consuming and incompatible with the alert rate. Secondly, the detection is still incomplete (i.e., attacks are still missed by IDS). Improving the detection rate requires the proliferation of heterogeneous sensors, so as to enhance the monitoring coverage and benefit from complementary detection techniques. However, multiplying sensors also multiplies the number of alerts received by security operators. There is a need for intrusion detection sensors to collaborate and exchange information.

2.2. Alarm correlation

Alarm correlation is a subfield of intrusion detection, whose goal is to make heterogeneous IDS sensors cooperate, in order to improve the attack detection rate, enrich the semantics of alerts and reduce the overall number of alerts.

Alarm correlation cannot be summarized to a single step in the analysis of alerts. Except Valeur et al. [3], who propose a correlation workflow intended to unify the various correlation steps, most correlation approaches proposed in the literature generally focus on specific aspects of the alert analysis.

These correlation approaches can basically be split in two categories, namely the *implicit* and *explicit* ones. Our objective here is not to review all alarm correlation approaches, but to briefly sketch some of them.

Implicit alarm correlation uses data-mining paradigms in order to fuse, aggregate and cluster large alert datasets. For example, the approaches of Valdes and Skinner [4], Dain and Cunningham [5,6], as well as Debar and Wespi [7] are based on the similarity between alert features (e.g., IP address of the victim and attacker). These processes are crucial to facilitate the analysis of the huge number of intrusion alerts, but generally fail to enhance the semantics of the alerts. Some extensions of these approaches have been proposed to extract relevant information from alert groups, for example by mining association rules between alerts [8]. In [2], Julisch proposes to apply attribute oriented induction techniques in order to generalize alert groups and support root cause analysis. In [9], we proposed an extension of this approach inspired by logical concept analysis, where alarm correlation is tackled as an information retrieval problem. Logical concept analysis unifies querying and navigation of information, which facilitates the investigation of large alert datasets.

Explicit alarm correlation approaches rely on a language which allows security experts to specify logical and temporal constraints between alert patterns in order to recognize complex attack scenarios, which generally require several steps to achieve their ultimate goal. When a complete or a partial intrusion scenario is detected, a higher level alert is generated. For example, in [10], we proposed an explicit correlation scheme based on the formalism of chronicles, and in [11,12] an imperative language to correlate sequences of alerts. Cuppens and Ortalo [13] also proposed a similar language.

An extension of explicit alarm correlation approaches, sometimes referred to as semi-explicit, uses the assumption that complex intrusion scenarios are likely to involve attacks whose prerequisites correspond to the consequences of some earlier ones [14–16]. Therefore, semi-explicit correlation consists in associating preconditions and postconditions, represented by first order formulas, with individual attacks or actions. The correlation process receives individual alerts and tries to build alert threads by matching the preconditions of some attacks with the postconditions of some prior ones.

2.3. Knowledge representation

Despite their differences, almost all of these correlation approaches share a common requirement: the availability of some knowledge about the characteristics of the attacks and the context in which they occur. However, the correlation approaches do not focus so much on how to *represent* the required knowledge, as to how to *use* this knowledge in their reasoning process. These alarm correlation paradigms have generally been implemented on an ad hoc basis and validated in specific environments or using proprietary formats. We claim that having a consistent data model is a prerequisite for any alert fusion and correlation techniques to be applied. This is why we focus on such model in this paper.

The knowledge representation problem has partially been addressed in previous papers, which we present in Section 9.

As a summary, our intent in this paper is not to propose a new alarm reasoning technique; our objective is to join together the atomic concepts and relations that are required to correlate alerts in a complete and consistent model, called M4D4 (=$(\text{M2D2})^2$). This model is to be used as a basis upon which existing and future cor-