Journal of Computer and System Sciences 80 (2014) 571–590

Contents lists available at SciVerse ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

Building a reputation-based bootstrapping mechanism for newcomers in collaborative alert systems

Manuel Gil Pérez [a,*], Félix Gómez Mármol [b], Gregorio Martínez Pérez [a], Antonio F. Skarmeta Gómez [a]

[a] Departamento de Ingeniería de la Información y las Comunicaciones, University of Murcia, 30071 Murcia, Spain
[b] NEC Europe Ltd., Kurfürsten-Anlage 36, 69115 Heidelberg, Germany

ARTICLE INFO

Article history:
Received 30 November 2012
Received in revised form 30 April 2013
Accepted 14 June 2013
Available online 4 July 2013

Keywords:
Reputation bootstrapping
Cold-start problem
Trust model
Collaborative detection networks
Distributed alert systems

ABSTRACT

Today trust is a key factor in distributed and collaborative environments aimed to model participating entities' behavior, and to foresee their further actions. Yet, prior to the first interaction of a newcomer in the system trust and reputation models face a great challenge: how to assign an accurate initial reputation to a newcomer? The answer needs to tackle two well-known problems: *cold-start* and *reputation bootstrapping*. Cold-start is a common issue to any system when newcomers boot for the first time, while reputation bootstrapping especially affects highly distributed scenarios, where mobile entities travel across domains and collaborate with a number of them. In this paper we focus on the two problems, which are addressed through a novel reputation bootstrapping mechanism for newcomers in a collaborative alert system aimed at detecting distributed threats. Experiments confirm the accuracy of our proposal as well as its robustness in the presence of ill-intentioned entities.

© 2013 Elsevier Inc. All rights reserved.

1. Introduction

Monitoring and analysis capabilities of existing Intrusion Detection Systems (IDS) are limited to a small portion of the overall network to which they belong. In this sense, the alarms produced by them are considered as isolated incidents, as they would make little sense if analyzed individually [1]. For example, attackers can launch a UDP port denial-of-service attack, described in [2] by the Computer Emergency Response Team, by injecting a large number of dummy packets over the entire network. They can exploit this vulnerability without being detected by overcharging services placed at different network segments (e.g., by sending out ICMP ECHO packets). The IDSs installed in each network segment may not trigger alarms if the number of packets sent to each segment is below a certain threshold. Nevertheless, the entire network could become congested, and thus compromised.

1.1. Collaborative alert systems

Building a collective knowledge base of alerts is a prerequisite for detecting distributed threats, such as the one introduced above. This can be acquired by sharing the alerts that each IDS produces individually, thereby achieving a more global perception about what is happening in the entire network. Therefore, IDSs need to establish a *Collaborative Alert System* (CAS) capable of enhancing the accuracy in detecting distributed threats [3].

* Corresponding author. Fax: +34 868 884151.
E-mail addresses: mgilperez@um.es (M. Gil Pérez), felix.gomez-marmol@neclab.eu (F. Gómez Mármol), gregorio@um.es (G. Martínez Pérez), skarmeta@um.es (A.F. Skarmeta Gómez).

0022-0000/$ – see front matter © 2013 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.jcss.2013.06.012

In a traditional network, IDSs are installed, configured, and updated by system administrators who possess the necessary technical skills to do that as experts in security [4,5]. We name these types of entities as *static IDSs*, as they are permanently placed in a given zone of the network to report alerts about suspicious events in their monitoring scope. In despite of that, the static configuration of networks is evolving towards a new paradigm where mobility of users is hoarding a great interest. This fact is prompted by users who can own devices at affordable prices (laptops and smart phones) with high computing performance and new added-value features, such as wireless communication capabilities. In the context of collaborative alert systems, organizations can leverage this opportunity to gather further security alerts that roaming users' devices can produce while temporarily staying in their domains.

Roaming users can (voluntarily) join to collaborate with security domains by sharing the alerts that they are capable of producing. In order to encourage cooperation, involved organizations can reward roaming users with some incentives, for example by granting them higher bandwidth. We name these types of entities as *mobile IDSs*, as they present special capabilities in mobility.

Regardless of the type of entity, the CAS has to be sure that all its IDSs exhibit a good and expected behavior about the alerts that they share. Disclosing bogus alerts can compromise the system to reduce its accuracy in detecting distributed threats. They will usually come from malicious or misbehaving IDSs, possibly due to an internal malfunctioning or as a consequence of being compromised by an attacker. Thus, the publication of bogus alerts from malicious IDSs can mistakenly lead the CAS to believe that a threat has actually occurred, when it really has not.

1.2. Motivation

A great number of trust and reputation models have emerged in recent years to measure the *goodness* of information sources [6,7]. They provide a way of finding out bad behaviors (alerts shared by malicious IDSs, in our context) by analyzing all interactions they have had with the system; either interactions with a given domain (*direct experiences*) or by requesting recommendations to others (*indirect experiences*) [8]. In any case, almost all of these models assume that the system is up and running, where all entities have a certain reputation score. This assumption is, however, hard enough to realize in practice due to

- at the beginning, no entity has yet interacted with any other (there is no historical information about anyone);
- a new entity can join, or can be connected to, the system for the very first time (there is no information about this newcomer); or
- in highly distributed networks, mobile entities can travel across heterogeneous domains with which they have not yet interacted.

The first two cases are often known in recommender systems as the *cold-start* problem [9], while the latter is related to the *reputation bootstrapping* problem [10]. Both problems have been widely discussed in the literature [11,12] although, as far as we know, this is an unexplored research area in the context of collaborative alert systems. A static or mobile IDS, or any other entity that wants to collaborate with the CAS, will be involved in the cold-start problem at the least once, when interacting with the system for the first time. Nevertheless, only mobile IDSs can be involved in the reputation bootstrapping problem as they are the only ones that can join and leave domains freely at any time.

1.3. Our contribution

We propose a novel reputation-based bootstrapping mechanism aimed at assigning an initial reputation score to a newcomer, either a static or mobile IDS, to encompass the following two situations: no historical information is available about the newcomer (cold-start problem) and, in highly distributed networks, historical information is available but in other domains where the newcomer has collaborated (reputation bootstrapping problem). We also stretch this mechanism to compute the initial trust that a security domain can deposit on other domains when they want to establish a new trust relationship each other.

In this paper, we have identified the three main actors (newcomers) that can join a security domain for collaboration purposes, namely:

- *static IDSs*, permanently placed in a network to monitor local services and resources of a specific network zone;
- *mobile IDSs*, belonging to users who want to collaborate with security domains to gain certain benefits; and
- *security domains* that want to improve their accuracy in detecting distributed threats by exchanging security information (mainly alerts) with other trusted domains.

Capabilities, or *detection skills*, that an IDS can offer to a security domain are used in this paper to compute its initial reputation score. In this sense, a security domain can assign a higher initial reputation score to whom can supply alerts related to threats uncovered by the rest of the domain's IDSs, or even regarding certain detection zones covered by "suspicious" IDSs (those with a low reputation score). On the other hand, computing the initial trust between two security domains also relies on their detection skills, by quantifying their *similarity* to detect common threats in which they are

Download English Version:

https://daneshyari.com/en/article/429815

Download Persian Version:

https://daneshyari.com/article/429815

Daneshyari.com