Expert Systems With Applications 102 (2018) 12–26

Contents lists available at ScienceDirect

Expert Systems With Applications

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

A self-organising multi-agent system for decentralised forensic investigations

Phillip Kendrick[a,*], Natalia Criado[b], Abir Hussain[a], Martin Randles[a]

[a] John Moores University, Liverpool, L3 3AF, United Kingdom
[b] King's College, London, WC2R 2LS, United Kingdom

ARTICLE INFO

Article history:
Received 6 August 2017
Revised 14 February 2018
Accepted 15 February 2018
Available online 15 February 2018

Keywords:
Multi-agent systems
Cyber security
Network forensics

ABSTRACT

As network-based threats continue to evolve more rapidly, detecting and responding to intrusion attempts in real-time requires an increasingly automated and intelligent response. This paper provides an agent-based framework for the analysis of cyber events within networks of varying sizes to detect complex multi-stage attacks. Agents are used as intelligent systems to explore domain specific and situational information showing the benefit of adaptive technologies that proactively analyse security events in real time. We introduce several algorithms to encapsulate and manage the traditional detection technologies and provide agent-based performance introspection as a mechanism to identify poorly performing systems. Our evaluation shows that the algorithms can reduce the amount of processing needed to analyse a security event by over 50% and improve the detection rate by up to 20% by introducing corrective systems to reduce false alarm rates in error-prone environments.

1. Introduction

With the increasing size of networks and the requirement for organisations to share business-critical information, current cyber security solutions, such as Intrusion Detection Systems (IDS) (Mukherjee, Heberlein, & Levitt, 1994; Verwoerd & Hunt, 2002) and manual network forensics (Clint, Reith, Carr, & Gunsch, 2002), have been unable to adapt to modern requirements. The increasing use of mobile and wireless technologies has expanded the boundaries of the traditional network by introducing a dynamic component wherein users and devices may come and go as needed. In addition to this, the pervasive adoption of the Software As A Service paradigm, characteristic of cloud-based software that can be updated or changed with ease, can alter the network's shape by enabling or disabling services and protocols. Furthermore, specific structures such as supply chain networks can increase the digital attack surface if not well protected by scalable security models (Zolfpour-Arokhlo, Selamat, & Hashim, 2013). Within this context, traditional security technologies have been unable to scale to the necessary levels due to their centralised nature and expensive hardware limiting their application to a single fixed network model (Liao, Richard Lin, Lin, & Tung, 2012).

IDSs are most commonly deployed as either network-based IDSs (NIDS) or host-based IDSs (HIDS). The network-based variation has access to raw packet data (Mahoney & Chan, 2001) collected directly off the wire which provides insight into how the endpoints[1] within the network communicate with each other, while NetFlow data (Galtsev & Sukhov, 2011) consists of aggregated statistics about the packet data. The host-based variant is installed on individual machines and has access to user-specific data such as the contents of decrypted packets and biometric data (e.g., keyboard typing speed and system calls) (Rudrapal, Das, Debbarma, & Debbarma, 2013). IDSs can further be categorised into signature-based, misuse-based and anomaly-based detection (Carvalho, Barbon, Mendes, & Proença, 2016) depending upon the model used for analysis. Signature-based detection uses a database of predefined examples of malicious activity to identify attacks. Signatures are defined by domain experts after a new attack has been detected to match all future instances of that attack. The manual process of defining signatures prevents the detection of previously unseen zero-day attacks making them a reactive technology. Misuse and anomaly-based detection provide an alternative by attempting to detect deviations from normal behaviour, as in the case of anomaly detection, or by learning the characteristics of abnormal behaviours and pattern matching future instances

* Corresponding author.
E-mail addresses: p.g.kendrick@2012.ljmu.ac.uk (P. Kendrick), natalia.criado_Pacheco@kcl.ac.uk (N. Criado), a.hussain@ljmu.ac.uk (A. Hussain), m.j.randles@ljmu.ac.uk (M. Randles).

[1] An endpoint is defined as any networked device within the internal network that has an IP address and is capable of communication; examples include computers, mobile devices and networked services.

https://doi.org/10.1016/j.eswa.2018.02.023
0957-4174/

as in the case of misuse detection (Tsai, Hsu, Lin, & Lin, 2009). Misuse and anomaly-based detection are seen as more flexible and scalable than signature-based approaches but may result in a higher number of incorrectly classified instances due to the lack of grounded knowledge (i.e., known signatures) (Zuech, Khoshgoftaar, & Wald, 2015).

In this paper, we propose a Decentralised Multi-Agent Security System (DMASS) as a scalable solution for the collection and analysis of cyber security and network forensic data (Kendrick, Hussain, & Natalia, 2016). The proposed DMASS model adapts to changing network architectures through the introduction of new agents and is far more scalable than current IDS solutions, which will typically require expensive high-end hardware to avoid performance bottlenecks (Verwoerd & Hunt, 2002).

The proposed multi-agent approach uses a collection of agents, which are distinguished from traditional software by their autonomous implementation, to perform a variety of roles in the network security environment. In addition to performing network monitoring and attack detection, currently carried out by IDSs, this research focuses on bestowing agents with the tools to replicate the manual forensic process, currently conducted by trained practitioners, to examining the security environment pragmatically. Bestowing agents with the ability to react to environmental changes, consider the performance of other agents and to work proactively to follow one line of investigation over another, when there is evidence to support it, is the fundamental principle included in the proposed model. This approach to digital evidence collection and cyber security is different from the traditional IDS approaches that typically use either signature, anomaly or misuse detection (Zuech et al., 2015). The DMASS approach of using automated forensic processes increases the agent's adaptability by enabling it to respond to unforeseen circumstances where the attacker can evade traditional signature or anomaly detection.

The remainder of this paper is organised as follows. Section 2 contains an analysis of the current research in the area. Section 3 contains an outline of our DMASS model. Section 4 contains a case study to illustrate the benefits of using our approach. Section 5 formalises the concept of domains by describing how information is obtained by the agents. Section 6 describes the agent simulator developed for testing the proposed systems. Section 7 describes algorithms to aggregate agent decisions and information collected by agents to improve the efficiency and detection performance. Finally, Section 8 contains conclusions and a discussion on future work.

2. Related research

In this section, we review existing agent-based architectures used for assessing network security. Ideally, agents should take advantage of the scalability and deployability improvements offered by multi-agent architectures and avoid common problems experienced with centralised processing, expensive hardware and rigid operating structures.

Shakarian, Simari, Moores, and Parsons (2015) use an agent-based cyber attribution system with agent reasoning to consider multiple sources of information. Agents use information derived from various military sources to reason about factors such as the geographical location, political landscape and possible motives of an attack. This system has the advantage of using high-quality information sources which are used to make conclusions about cyber attacks. The system classifies data as either a fact or presumption, treating presumptions as unverified facts. The DMASS architecture presented in this paper is designed for use in non-military fields and recognises that data may be incorrect or missing requiring strategies to seek out information proactively. Furthermore, agents perform live data collection to gather the most up-to-date information available to avoid problems of data degradation often experienced with central information repositories.

Haack et al. (2010) use a hierarchical Multi-Agent System (MAS) for monitoring and reporting policy violations within the security environment. The system is composed of various agent types each with a particular task to perform (e.g., event monitor, alert generator, report builder). The network administrator defines a network policy for the agents to implement in a hierarchical model with instructions passed down from higher to lower-tier agents. This model is inherently centralised and suffers from many highlighted disadvantages experienced by IDSs. The top-layer agents classify security events[2] from their fixed position using data collected by the mobile lower-tier agents. The layered structure of having one class of agents to collect and another class of agents to analyse produces a static information flow which can suffer from availability downtime and performance bottlenecks. A more scalable approach, advocated in this paper, is to allow each agent to make decisions about the security events from their local viewpoints which are then brought together to classify the event as a whole.

Jahanbin, Ghafarian, Seno, and Nikookar (2013) introduce an agent framework for forensic information gathering using three types of agents for data collection, analysis and alert generation. The authors remark that the MAS paradigm is well suited to the task of forensic data collection since agents can be dispatched to areas of the network to perform evidence gathering, a feature lacking in many IDSs that just monitor the visible network connections. Structural similarities exist with the system proposed by Haack et al. (2010), i.e., with three layers of agents forming an information pipeline from the lower layers to a higher layer agent. The decision-making process used in this model is similar to an IDS because the security decisions are made based upon the available data without consideration of possible missing data. Our proposed DMASS evaluates security events based on what data is found as well as what is missing; accounting for the possibility that the attacker may have obfuscated evidence during the attack.

Shanmugasundaram, Memon, Savant, and Bronnimann (2003) develop a distributed forensics system using a hierarchical approach with multiple configurable sensors placed on the network. The system uses a variety of sensors and servers to collect and aggregate the information to derive the nature of the security event from the observable data. The system identifies the attack type based on which pieces of evidence are missing during the search. In the complex and changing cyber security environment, this approach is desirable since the lack of information does not necessarily conclude no attack has taken place.

Baig (2012) survey the current applications of MASs in critical infrastructure fields including intrusion detection. System resilience is highlighted as an important factor for multi-agent architectures where attackers could force agents offline through denial of service attacks. Agents should be able to adapt to changes in the network structure by taking into account agents which may not be able to communicate or devices that cannot be reached. Hierarchical models cannot function unless their communications pipeline is unimpaired which is unrealistic in modern expanded networks. Our proposed system uses heterogeneous agents that operate without a central control structure to withstand attacks on the availability of the system. To further improve the resilience of agents, communication paths and the performance of agents are determined at run-time based on the attack characteristics, for example, to avoid communicating over areas of the network currently under attack.

Mees (2012) uses a MAS to detect Advanced Persistent Threats (APTs) (Chen, Desmet, & Huygens, 2014) by using external data

[2] A security event is defined as one or more series of suspected attacks.