Available online at www.sciencedirect.com
ScienceDirect
Computer Networks 51 (2007) 632–654
Computer Networks
www.elsevier.com/locate/comnet

Improving the quality of alerts and predicting intruder's next goal with Hidden Colored Petri-Net ☆

Dong Yu *, Deborah Frincke

Department of Computer Science, University of Idaho, Moscow, ID 83844-1010, United States

Received 3 August 2005; received in revised form 27 January 2006; accepted 19 May 2006
Available online 22 June 2006

Responsible Editor: Christos Douligeris

Abstract

Intrusion detection systems (IDS) often provide poor quality alerts, which are insufficient to support rapid identification of ongoing attacks or predict an intruder's next likely goal. In this paper, we propose a novel approach to alert postprocessing and correlation, the Hidden Colored Petri-Net (HCPN). Different from most other alert correlation methods, our approach treats the alert correlation problem as an inference problem rather than a filter problem. Our approach assumes that the intruder's actions are unknown to the IDS and can be inferred only from the alerts generated by the IDS sensors. HCPN can describe the relationship between different steps carried out by intruders, model observations (alerts) and transitions (actions) separately, and associate each token element (system state) with a probability (or confidence). The model is an extension to Colored Petri-Net (CPN). It is so called "hidden" because the transitions (actions) are not directly observable but can be inferred by looking through the observations (alerts). These features make HCPN especially suitable for discovering intruders' actions from their partial observations (alerts) and predicting intruders' next goal. Our experiments on DARPA evaluation datasets and the attack scenarios from the Grand Challenge Problem (GCP) show that HCPN has promise as a way to reducing false positives and negatives, predicting intruder's next possible action, uncovering intruders' intrusion strategies after the attack scenario has happened, and providing confidence scores.

Keywords: Intrusion detection; Alert correlation; Hidden Colored Petri-Net

1. Introduction

One of the most important requirements of a good intrusion detection system (IDS) is the generation of high quality alerts. Unfortunately, IDS sensors usually generate massive amount of alerts [1], especially if those sensors have high sensitivity to potential misuse, as can be the case with tightly tuned anomaly based sensors. In the distributed

☆ This is an extended and enhanced version of the work published in the International Conference on Applied Cryptography and Network Security, Yellow Mountain, China, 2004 [39].

* Corresponding author. Tel.: +1 425 707 9282; fax: +1 425 706 7329.
E-mail addresses: dongyu@csds.uidaho.edu (D. Yu), frincke@cs.uidaho.edu (D. Frincke).

1389-1286/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2006.05.008

case, the situation is compounded because there are more sensors (and hence more data), and greater chance of time delay before sensor alerts are consolidated – with commensurate risk that some information will be inaccurate or stale. To make things worse, only 1% of the enormous amount of alerts generated by most IDS corresponds to unique attacks [1–3]. The remainders are false positives (i.e., alerts on non-intrusive actions), repeated warnings for the same attack, or alert notifications arising from erroneous activity or configuration artifacts.

Many ways are available to improve the quality of alerts. For example, we may achieve the goal by using better sensors, signatures, or analysis algorithms. In this paper, we focus on algorithms, and propose a novel alert correlation approach. In other words, our approach aims to reduce the false positives and false negatives by postprocessing (i.e., correlating) the alerts in a novel way: inferring an intruder's actions with a model named Hidden Colored Petri-Net (HCPN). The architecture of our system is depicted in Fig. 1. Raw audit data collected by sensors are first analyzed to generate alerts. These alerts, which contain large amounts of false positives, are then fed into our HCPN-based alert correlators to remove most of false positives and repetitions. Confidence scores of alerts from the HCPN components installed at different sites are then fused with our extended Dempster-Shafer theory of evidence to further improve the quality of alerts before they are sent to the active responser to make the reaction decision.

Alert correlation [1,4–20] is used to (a) reduce the number of alerts that an IDS would generate to more manageable levels while still retaining strong detection capacities, (b) improve IDS correctness by reducing the false positives and negatives in the alerts generated by the IDS sensors, and (c) unveil an intruder's intrusion strategy after the attack has happened.

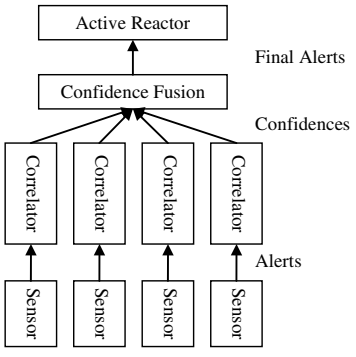

Fig. 1. Raw audit data collected by sensors are first analyzed to generate alerts. These alerts, which contain a large amount of false positives, are then fed into the alert correlators. Confidence scores provided by the alert correlators located at different sites are fused before being sent to the active responser to make the reaction decision.

One way to look at the alert correlation problem is to extract "true" alerts (or filter out the false alerts) from the raw alerts generated by the IDS sensors by utilizing relationships (e.g., similarities, sequential relationships, etc.) among alerts. This filter view of alert correlation, which will be discussed further in Section 5, has been taken by Cuppens et al. [13,14] and Ning et al. [17–19], to name a few. Approaches based on this filter view can remove (or filter out) large percentage of false positives. However, this formulation of the problem works directly upon the alerts. It does not distinguish between alerts and intruders' actions in the correlation process, and usually does not use information such as false negative rate and false positive rate to improve the correlation results.

In this paper, we take a different view and consider alert correlation as the problem of inferring an intruder's actions based on partial observations – alerts, progressively. Based on this perspective, we propose a novel methodology for analyzing alerts. Our approach is based on a theoretical model named Hidden Colored Petri-Net (HCPN). Based on this inference view, we assume an intruder's (either an actual intruder or an insider who is misusing the system) actions are unknown to the IDS and can be inferred only from the enormous amount of low quality alerts generated by the IDS sensors. We demonstrate that HCPN, as an extension to Colored Petri-Net [22,33,34], can describe the relationship between different steps carried out by intruders, model observations (alerts) and transitions (actions) separately, and associate each token element (system state) with a probability (or confidence). The model is called "hidden" because the transitions (actions) are modeled as hidden variables (i.e., not directly observable by the analyzer) in HCPN. However, it can look through the observations (alerts) to infer the transitions (actions). When no training data are available, HCPN behaves in a similar way as other alert correlation approaches that solely depend on the precondition–postcondition relationship. When training data are available, it can learn from the data and further improve the alert correlation result.