



Multilevel Fusion-Based Intrusion Detection

Remco de Boer

SemLab, Wassenaarseweg 72, 2333 AL Leiden, The Netherlands, deboer@semmlab.nl

Jan van den Berg

Erasmus University Rotterdam, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands, jvandenbergh@few.eur.nl

Wilco van Ginkel

Ubizen, Rendementsweg 20-B1, 3640 AD Mijdrecht, The Netherlands, wilco.vanginkel@ubizen.com

ABSTRACT

Shortcomings of current intrusion detection systems, most notably high false alarm rates and insufficient attack detection accuracy, call for a structured, sophisticated approach. We identify multi-sensor data fusion as such an approach and present a multilevel intrusion detection system architecture. At each level, logically independent functional units combine the data or information from various sources using the technique of data fusion. In this way, each unit contributes to the overall quality of the intrusion detection system. We present the set of functional tasks to be performed, their hierarchical relationships, and sketch the way the units should work together. The corresponding multilevel 'blackboard' architecture can be used as starting point for implementing next generation high quality intrusion detection systems.

INTRODUCTION

A common approach for dealing with information security is to install - in addition to preventative techniques related to identification, authentication, and authorization - systems for intrusion detection. Such intrusion detection systems (IDSs) act as a second line of defence [2] and are supposed to automatically offer intrusion detection functionality. Traditionally, IDSs are viewed as consisting of three, logically distinct, functional components [3]: the sensor, the analyzer, and the user interface. The sensor is responsible for collecting the data and the analyzer for determining whether an intrusion has occurred. The user interface enables an (usually human) IDS expert to inspect the output of the analyzer and to control the behaviour of the system.

The classification of the analyzer puts the IDS in one of two states: positive, indicating an intrusion, or negative, indicating no intrusion. Since the analyzer's conclusion is either correct ('true') or incorrect ('false'), there are a total of four classifications of the state of the IDS: true positive, true negative, false positive, and false negative. Ideally, false positives and false negatives never occur, true negatives are neglected completely, and true positives result in a correcting (automated or human) response, the strength of which depends on (a) the applied security objectives and (b) the gravity of the intrusion.

Regarding the detection mechanism of the analyzer, a distinction is made between two fundamentally different approaches [2]: misuse detection and anomaly detection. Misuse detection is based on a comparison of the observed data to a list of 'signatures' of known attacks where matches are reported as intrusions. Anomaly detection relies on the assumption that all intrusive actions are anomalous, i.e., anything that deviates from 'normal activity' is deemed an intrusion.

An elaborated analysis of current approaches for both anomaly detection and misuse detection [2] has made apparent that no single approach can detect all types of intrusions. Two fundamental problems concerning current IDSs are often put forward [3-5]: (a) high rate of false positive alarms, and (b) low attack detection accuracy. Actually, all current IDSs suffer from these shortcomings. In an attempt to remove them we shall apply a recent idea in the area of intrusion detection,

namely, multi-sensor data fusion [5]. The corresponding model was derived using a translation from an existing data fusion model for military purposes to a fusion model applicable to the area of intrusion detection. However, the model proposed needs still to be elaborated. In this article we present a structured approach for implementing such a multilevel fusion-based intrusion detection system. At each level of the corresponding architecture, logically independent functional units called 'experts' combine the data or information from various sources using the technique of data fusion. In this way, each expert contributes to the overall quality of the intrusion detection system.

In the remainder of this article, we start by recapitulating the fundamentals of multisensor data fusion. Next, we present the goal, multilevel decomposition, and some architectural issues of multilevel fusion-based IDSs. We finalize by giving some conclusions.

MULTISENSOR DATA FUSION

Current data fusion concerns a rapidly evolving engineering discipline [6-8]. Originating in the military domain in the late 1970s, data fusion methods in recent years have also been applied to problems in the civilian domain [7]. A biologically motivated fusion process model has also been developed [9].


Data fusion concerns a set of means and tools where data originating from different sources are combined in order to obtain information of 'greater quality' [10]. In practice, data fusion often applies a multilevel approach with several levels of abstraction. At each level, an appropriate data fusion technique has to be chosen. For example, in a centralized approach applied at the lowest level, either raw data may be fused directly ('data level fusion'), or feature vectors may be fused after feature extraction ('feature level fusion'), or individual estimates of the sensors may be fused ('decision level fusion') [7]. A decentralized approach with multiple experts (typically applied at higher levels) usually involves more complex data fusion processes. 'Cooperative fusion', for example, involves communication and cooperation between the various experts in order to interpret the global state of the system and to decide on this. The problem solving approach probably most used in higher level data fusion applications is 'blackboard processing' [11]. All independent knowledge sources (experts) have access to a central blackboard. All communication and interaction between the experts takes place using this blackboard, which is a global database containing the solution state. Each expert can use, change and create solution state data stored on the blackboard.

MULTILEVEL FUSION-BASED INTRUSION DETECTION SYSTEMS

Goal

Based on the sketch given in section 1, we argue that next generation IDSs should provide information of greater quality consisting of:

Fig. 1. Situational view of intrusions showing both attacks and attackers.



- *situational views* (instead of huge numbers of alarms);
- *minimized number of false positives*.

The alarms that result from current IDSs are in general too fine-grained and too low-level. Instead of presenting large numbers of alarms to the human network operators, high-level situational descriptions should be offered, i.e., courser views where both *attacks* and *attackers* are clearly identified. As an example, figure 1 depicts such a *situational view*.

It shows three different attacks, launched against three networks or network segments by three individual attackers. Attacker *x* is involved in a single attack against Networks 1 and 2. Attacker *z* is involved in a single attack against Network 3, but this is a coordinated attack in which Attacker *y* is also involved. Furthermore, Attacker *y* has also launched his 'own' attack against Network 1. Some actions that the attacker performs might be legitimate actions. Note also that not all of an attacker's actions - either legitimate or intrusive - are directly observable in the network or network segment. Situational views like the one presented above provide a *coarser* view on the occurrences of intrusions than 'traditional' IDSs. Therefore, they reduce the high number of positives that result from an intrusion. However, automatic generation of situational views is not a sinecure. We need a functional decomposition in several layers.

Multilevel Functional Decomposition

Starting point of our analysis is the functional decomposition. Inspired by the work of Tim Bass [5], five levels are chosen in the IDS Data Fusion Model.

Level 0: From Events to Alerts

At level 0, initial signal processing takes place. Host-based and network-based sensors are supposed to observe all kinds of low-level *events* typically being systems calls, application calls, and passing network packets. These events may be either just regular events or occurrences related to an intrusion. Therefore, filtering is necessary. This filtering can be based on the earlier mentioned techniques of misuse and anomaly detection. Events that pass the filters are called an *alert* and point to a *possible attack*.

Level 1: From Alerts to Alert Tracks

At level 1, data fusion is applied in order to find groups of alerts called *alert tracks* that are the result of the corresponding 'potential intrusion process(es)'. To do so, the alerts as resulting from the level 0 filter are first 'aligned' to a common reference frame. This alignment facilitates comparison and further processing of multiple alerts that originate from different sources. Next, during the 'association' stage, each alert is analyzed in order to decide whether it will be added to an existing alert track or whether a new alert track should be started.

Level 2: From Alert Tracks to Situational Views

The goal of the level 2 analysis is to identify two types of alert aggregations, namely alerts that together make up an attack and alerts that together represent the behavior of a single attacker. The aggregation results are put together in a set of 'situational views' as introduced and discussed in subsection 3.1.

Level 3: Prioritization of situational views

Through further reasoning on the situational view result, level 3 processing enables the determination of the threat of the current situation and, based on that, of the expected exposure. This evaluation takes the vulnerability of the target into account. Exposure assessment enables prioritization of attacks taking place, thereby aiding the user to focus on the most threatening attacks: the corresponding situational views are shown to the human IDS manager together with a set of intervention tools.

Level 4: Resource Management

Level 4 processing performs evaluation and feedback of the total fusion process. This concerns a meta-process for managing the proper working of the complete IDS: after each evaluation, it adjusts the processes at each level and cues the sensors to dynamically improve the quality of the output offered to the IDS manager.

Towards an Architecture

An architecture for fusion-based IDSs can be based on the multilevel functional decomposition presented above. We confine ourselves here to sketch implementations of the four lowest levels. A visualization of the corresponding architecture is given in figure 2.

From Events to Alerts: implementation issues

At the lowest level, filtering of events should be implemented to find alerts. Traditional IDSs can be used here. The underlying techniques are, as mentioned in the introduction, misuse detection and anomaly detection. It is the responsibility of the Resource Management process (level 4) to collect enough information for improvements of the filtering processes. For instance, in case an unknown series of events occurs, this set of suspicious events is filtered out and labeled as alerts. If later on (at higher levels), these alerts appear to be related to just regular activities not related to any kind of attack, the filter characteristics are adapted through the Resource Management process such that next time, similar series of events are simply ignored. The other way around is also possible, namely, in case an intrusion has taken place while the corresponding alerts had not been filtered out. This leads to a false negative. If afterwards the intrusion is discovered based on the damage done, inspection of the corresponding log-files may result into enforcement of the filtering processes at level 0.

From Alerts to Alert Tracks: implementation issues

The implementation of the level 1 functionality involves various activities. First of all, the level 0 alerts coming from different sources like network-based IDSs and host-based IDSs should be aligned into a standardized format. As common reference frame, the Intrusion Detection Message Exchange Format (IDMEF) data model [12] could be used. The alignment procedure concerns more than just formatting, it also relates to feature selection: only relevant features of events are collected and put together into a frame. Examples of relevant features are IP-addresses, message numbers, protocol id's, time stamps, id's of system calls, special characters, and so on.

The next step concerns *data fusion*. Specific combinations of level 0 alerts constitute together one alert track. New incoming alerts should be classified: if an incoming event is considered being related to an existing alert track, it may be added to that one. Otherwise, a new alert track should be started. Every time a new alert is assigned to an existing alert track, the confidence in the importance of the alerts in this track increases. Based on confidence characteristics of the sensor that reported an alert, the initial confidence in a single-alert track can be calculated. Furthermore, it seems inevitable to re-evaluate from time to time whether alerts have been assigned to the right alert track. To

Fig. 2. Blackboard architecture for a fusion-based IDS

implement the composition of the right alert tracks, intelligent techniques from the area of (statistical) pattern matching like neural networks and fuzzy systems need to be used. Decision trees and other machine learning algorithms might be helpful too.

The final step in level 1 processing is refining the estimate of the identity of the observed event [13]. If an event is reported for the first time it may be difficult to unambiguously determine its identity. Subsequent reports of the same event may contain data

that is supplementary to the information already known. This information is combined with that reported earlier. Hence, with every reported alert, the estimate of the identity of the event can be refined.

The level 4 management process should also interfere at level 1. Removal of alert tracks not being referred to during a certain amount of time is such an activity.

Towards Situational Views: implementation issues

The creation of situational views based on a set of alert tracks involves two different fusion processes since (a) a single attacker can be involved in multiple attacks and (b) multiple attackers can be involved in a single attack. Furthermore, not all attacker behavior is necessarily part of an attack while, on the other hand, all attacks are necessarily the result from attacker behavior. Finding alerts that represent the behavior of a single attacker can be viewed as a tracking problem: the actions of an attacker leave a 'trail of alerts' that must be combined (fused) by an 'Attacker tracker' process. At the same time an 'Attack recognizer' process should take care of the recognition of attacks. Background knowledge and sophisticated reasoning techniques [13] are needed for the implementation of these recognition processes. In our opinion, the ability of next generation IDSs to offer situational views should be considered as the most important improvement.

It should also be clear that the above-mentioned processes could make use of the same kind of (dynamically changing) information like alert tracks and all kinds of details about attacks and attackers. To enable a flexible cooperation of these processes, we propose to apply the so-called blackboard architecture [11] where common information is made available at a 'blackboard' which can be read and write by authorized entities.

In addition, the notion and concept of false positives becomes important at this level. At the lower levels there were no false positives since all alerts are direct representations of occurring events. From a level 2 point of view, however, an alert that is neither part of an attack nor the result of attacker behavior is a false alarm, which should simply be neglected.

Threat Assessment: implementation issues

Level 3 of our multilevel model concerns the implementation of threat assessment. Here, we can again make use of the above-introduced blackboard architecture by defining a set of appropriate processes that communicate using the blackboard. There are several solutions imaginable. Using environmental background information an 'Intent predictor' process might be introduced to estimate the intention of an attack(er), a 'Vulnerability assessor' process to judge the actual vulnerability, an 'Outcome predictor' for estimating the final outcome of an attack, and a 'Threat assessor' for assessing the corresponding menace and its impact. Based on the information coming out the latter process, the most threatening attacks and attackers can be determined and shown to the human intrusion detection manager.

For a visualization of some of these details, we once again refer to figure 2. Much more details on the realization of the proposed architecture including an evaluation of it, are available in reference [1].

CONCLUSIONS

In this paper, an elaboration on a general model and the outline of an architecture for multilevel fusion-based IDSs has been sketched. It has been made credible how an IDS having this architecture, may be able to overcome the two most fundamental problems of the current generation IDSs namely high false alarm rates and low attack detection accuracy. At the moment, great efforts are going on to implement parts of the proposed architecture starting by implementing the level 1 recognition processes. Eventually, this implementation should yield proof of concept.

ENDNOTES

1 This paper may be considered as a condensed version of master thesis [1].

REFERENCES

- [1] de Boer, R.: A Generic Architecture for Fusion-Based Intrusion Detection Systems. *Master Thesis*, Erasmus University Rotterdam, Rotterdam School of Economics (2002) (Available at <http://www.few.eur.nl/few/people/jvandenbergh/masters.htm>, #40)
- [2] Biermann, E., Cloete, E., Venter, L.: A Comparison of Intrusion Detection Systems. *Computers & Security* (2001), 676–683
- [3] Allen, J., et al.: State of the Practice of Intrusion Detection Technologies. *Technical Report*, Carnegie Mellon Software Engineering Institute (2000)
- [4] Briney, A.: New Directions in Intrusion Detection. *Information Security Magazine* (2001)
- [5] Bass, T.: Intrusion Detection Systems and Multi-sensor Data Fusion. *Communications of the ACM* **43** (4) (2000), 99–105
- [6] Waltz, E., Llinas, J.: *Multisensor Data Fusion*. Artech House, Inc. (1990)
- [7] Hall, D., J.Llinas: An Introduction to Multi-sensor Data Fusion. *Proceedings of the IEEE* **85** (1997)
- [8] Yannone, R.: Exploring Architectures and Algorithms for the 5 JDL/DFS Levels of fusion required for advanced fighter aircraft for the 21st century (1999) (Abstract available at <http://65.105.56.185/master23/category159/A276193.html>)
- [9] Antony, R.: *Principles of Data Fusion Automation*. Artech House, Inc. (1995)
- [10] Wald, L.: Definitions and Terms of Reference in Data Fusion. *International Archives of Photogrammetry and Remote Sensing* **32** (1999)
- [11] Nii, H.: Blackboard Systems: The Blackboard System of Problem Solving and the Evaluation of Blackboard Architectures. *The AI Magazine* **VII** (2) (1986), 38–53
- [12] Group, I.D.W.: Intrusion Detection Message Exchange Format (2002) (Available at <http://www.ietf.org/html.charters/idwg-charter.html>)
- [13] Hall, D.: *Mathematical Techniques in Multisensor Data Fusion*. Artech House, Inc. (1992)

0 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:

www.igi-global.com/proceeding-paper/multilevel-fusion-based-intrusion-detection/32377

Related Content

Robotics and Programming Integration as Cognitive-Learning Tools

Nikleia Eteokleous (2018). *Encyclopedia of Information Science and Technology, Fourth Edition* (pp. 6859-6871).

www.irma-international.org/chapter/robotics-and-programming-integration-as-cognitive-learning-tools/184382

Analyzing Evolution Patterns of Object-Oriented Metrics: A Case Study on Android Software

Ruchika Malhotra and Megha Khanna (2019). *International Journal of Rough Sets and Data Analysis* (pp. 49-66).

www.irma-international.org/article/analyzing-evolution-patterns-of-object-oriented-metrics/251901

Practically Applying the Technology Acceptance Model in Information Systems Research

Reza Mojtahed and Guo Chao Peng (2013). *Information Systems Research and Exploring Social Artifacts: Approaches and Methodologies* (pp. 58-80).

www.irma-international.org/chapter/practically-applying-technology-acceptance-model/70710

Bi-Directional Business/IT Alignment

Hashim Chunpir, Frederik Schulte, Yannick Bartens and Stefan D. Voß (2018). *Encyclopedia of Information Science and Technology, Fourth Edition* (pp. 601-608).

www.irma-international.org/chapter/bi-directional-businessit-alignment/183774

Design of a Structured Parsing Model for Corporate Bidding Documents Based on Bi-LSTM and Conditional Random Field (CRF)

Lijuan Zhang, Lijuan Chen, Shiyang Xu, Liangjun Bai, Jie Niu and Wanjie Wu (2023). *International Journal of Information Technologies and Systems Approach* (pp. 1-15).

www.irma-international.org/article/design-of-a-structured-parsing-model-for-corporate-bidding-documents-based-on-bi-lstm-and-conditional-random-field-crf/320645