2013 5th International Conference on Cyber Conflict
K. Podins, J. Stinissen, M. Maybaum (Eds.)
2013 © NATO CCD COE Publications, Tallinn

Architecture for Evaluating and Correlating NIDS in Real - World Networks

Robert Koch
Faculty of Computer Science
Universität der Bundeswehr München
85577 Neubiberg, Germany
robert.koch@unibw.de

Mario Golling
Faculty of Computer Science
Universität der Bundeswehr München
85577 Neubiberg, Germany
mario.golling@unibw.de

Abstract: Research in the field of IT security - in this case especially the Evaluation and Correlation of Intrusion Detection Systems (IDS) - implies special demands for the construction and operation of IT systems. In order to (i) evaluate multiple IDS under absolutely identical conditions and to (ii) check their reactions especially against novel attack patterns / attacker behaviour, all attack related actions (i.e. all traffic) have to be forwarded to all IDS in parallel at real-time. In addition, an attractive target needs to be offered to potential attackers, awaking the outward semblance of real-productive systems / networks including the corresponding behaviour.

In particular, the correlation of IDS seems a promising approach to compensate the individual deficiencies of IDS. For example, while knowledge based systems are only able to detect previously known attacks, anomaly based systems suffer from higher False Alarm Rates (FARs). Even more, periodic performance evaluation studies, e.g., by NSS-Labs, have illustrated that numerous IDS are not configured properly and have a much worse system performance and detection capability than announced by the vendors. However, changing parameters of systems in productive networks (for the correlation of IDS as well as for their evaluation) can result in an enhanced endangerment of the security or even a breakdown of the network in case of horrible misconfigurations.

To overcome these shortcomings, we present an architecture that supports research in the field of IT security and simultaneously ensures that all actions associated with an attack get recorded and a spill over of the attack from the research to the productive environment is prevented. Each test system is supplied with an unaltered live record of the network traffic. This allows an assessment of the detection as well as a comparison of different NIDS concepts/products. In addition, different correlation strategies of alerts of multiple systems can be evaluated. Furthermore, superior configurations can be identified and assessed without endangerment of the productive network.

Keywords: *intrusion detection, optimization, real-world evaluation, comparative evaluation, intrusion detection correlation, test environment*

1. INTRODUCTION

Detecting and Defending attacks against networks and systems is an intense research area for over 30 years. Although a high number of security mechanisms have been developed, for instance numerous proprietary as well as Open Source (Network) Intrusion Detection Systems (NIDS), the situation is not easier. In contrast, the number of attacks, security incidents and malicious software is increasing constantly during the last years. So does the quality of the attacks. Nowadays, attacks are much more targeted and technically sometimes very complex. Even more, attack toolkits which perform sophisticated automated attacks are available on the underground market and can be purchased with Service Level Agreements, guaranteeing that the product will not be detected by todays widely used IDS for a specific amount of time, resulting in multimillion Dollar losses caused by cyber-crime every year.

In order to evaluate the protection mechanisms of existing and newly developed IDS, exposing them to real word attacks seems beneficial. Thus, in order to perform analyses on the reactions of IDS on novel attack patterns / attack behaviours, a secured and controlled research environment - where real attacks are allowed knowingly - is almost indispensable.

Up to now, also modern IDS are often not able to detect sophisticated attacks [1]. This is not only because of new and yet unknown attack techniques, but also because of misconfigurations, erroneous detection engines, etc. For example, studies by NSS-Labs have shown that most systems are configured badly. In addition, it has been demonstrated, that the detection performance in real-word networks by current IDS can be much worse than specified by the producer [2]. E.g., one system was only able to analyse 3 percent of the expected amount of traffic. Besides that, depending on the classification of the IDS, also several shortcomings can be found. For example, while knowledge-based systems are only able to detect already known attacks, anomaly-based systems suffer from higher False Alarm Rates (FARs).

Summarized, the most important real-world problems regarding the use of state-of-the-art IDS are:

- High False Alarm Rates
- Undetectable attacks
- Complex configuration
- Intense administration
- Data Encryption

Therefore, an environment is required which enables a testing and optimization of parameters/configurations as well as an in-depth evaluation of the performance and detection capabilities of IDS without an endangerment of the productive network. Such an environment can also be used for research and development of, e.g., correlation techniques for IDS. Especially with the exchange of attack information between different installations respectively of different autonomous systems spread around the globe and the corresponding analysis/correlation, it is possible to defend against new attack waves. For example, the Internet Storm Center of the SANS Institute collects data from over 500.000 Sensors around the globe. ATLAS from Arbor Networks or the exchange of statistical data by Cisco IPS are other examples for collaboration and the generation of early warnings.

In order to (i) assist the consumers by selecting the corresponding NIDS that best fits into their environment and to (ii) support the combination and correlation of different NIDS (like a anomaly-based with a knowledge-based NIDS), a common environment is needed which provides the network operator with the ability to install more than one NIDS without an influence on the productive network on the one hand and among the different NIDS on the other hand.

Taking into account the idea of evaluating and correlating NIDS, a comprehensive architecture that allows secure and manageable research within a subnet of an - otherwise productively used - network is presented in this publication. Each isolated NIDS is provided with a copy of the sniffed data traffic. In contrast to evaluations of system performances and capabilities with the help of synthetic data-sets like Lincoln Lab DARPA intrusion detection evaluation 98/99 [3], our architecture supports real-world data and real-time detection capabilities for a realistic system assessment. An injection of malicious traffic onto the productive network by the NIDS is prevented and also no modification of traffic during transit is possible.

The operator of the network will have the opportunity to use the proposed architecture on three different modes:

- **Test environment:** The customer implements a productive IDS next to one or more test IDS environments of the same system. These test IDS environments will be used to modify and optimize IDS policies. The customers can verify and validate their modifications without risking harmful influence on the network.
- **Validation of different systems:** For those customers that don’t know which product fits their purposes best, the proposed architecture can be used to test different IDS. Thus the customer will be able to validate the IDS implementations against each other based on real traffic. The customers will use this mode in order to choose the best system they will later implement into their network.

- **Correlation of the results of different NIDS:** As IDS will often produce false-positive and false-negative alerts, it would be beneficial to have multiple IDS implemented in order to validate results and to achieve a common operational picture. For this operational picture the evaluation and correlation unit is placed outside the test environments.

The paper is structured as follows: In Section 2, a practical scenario of research on Intrusion Detection will be given. In Section 3, we will discuss related evaluation techniques and approaches as well as requirements for the architecture. Based on that, the architecture of the system will be presented in detail in Section 4. Finally in Section 5, we will present results of a proof-of-concept implementation as well as a case study to show the benefits of the architecture, before the conclusions are drawn.

2. SCENARIO

In this section, the need for a holistic architecture for evaluating and correlating NIDS is illustrated using a practical, real-world scenario (see Figure 1). The special feature of the scenario is the integrative approach of different components; from Intrusion Detection (through multiple sensors) over live analysis (automated correlation of data) to post-mortem analysis (IT forensics). In the following, the individual components are presented:

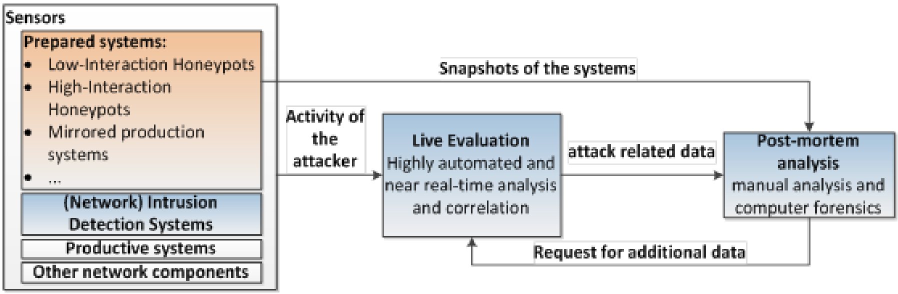
Sensors
Prepared systems:
Low-Interaction Honeypots
High-Interaction Honeypots
Mirrored production systems
...
(Network) Intrusion Detection Systems
Productive systems
Other network components
Snapshots of the systems
Activity of the attacker
Live Evaluation
Highly automated and near real-time analysis and correlation
attack related data
Post-mortem analysis
manual analysis and computer forensics
Request for additional data

Figure 1. Overview of the components

A. SENSORS

Attackers are attracted with specially prepared systems (clients running Windows XP, Windows 2003 servers as well as low-interaction and high-interaction honeypots) with deliberately (simulated) vulnerabilities in the research environment. The course and the behaviour of the attacker are recorded multiple times – both by host and network components (host / network intrusion detection systems, honeypots, switches with monitoring port, etc.).

B. LIVE EVALUATION

The sensors are forwarding the recorded data to a central database for analysis and correlation of the activities of the attackers. The alerts generated form the basis for further investigations. The evaluation of the high amount of alerts is first carried out by an automatic correlation. For this, already several approaches are existing, which - however - all have individual shortcomings and needed to be improved [4, 5,6]. Current approaches only consider alerts of IDS, but no additional sources such as Honeypots and log data included.

C. POST-MORTEM ANALYSIS

For the reconstruction of an attack, additional data or snapshots can be requested using specific criteria, such as time stamps, source or destination. This extends the database for forensic examinations. Finally, automated countermeasures can be taken through Intrusion Prevention Systems (IPS).

3. RELATED WORK

A. INTRUSION DETECTION SYSTEMS

IDS can be classified using numerous characteristics, where the most important one is the detection technique. Two different mechanisms are used, namely:

- **Knowledge-based detection**: Detection of attacks by the use of knowledge about malicious events, e.g., by matching a set of known misuse patterns (signatures) against a stream of packets or events.
    - ○ **pro:** low false alarm rates, precise diagnostics
    - ○ **con:** insensitivity to attack variations, difficulty of signature maintenance/ updates, incompleteness of known patterns, huge databases, difficult to reach near-real-time evaluation of network links with high data rates
- **Anomaly-based detection:** Detection of attacks by measuring deviation from statistical models of normality.

  - **pro:** detection of unknown attacks, adjustment to traffic/process drift
  - **con:** high false alarm rates, anomalies ≠ attacks

Both methods of detection have their pros and cons. Often, knowledge based systems are preferred since they typically provide lower FARs and precise diagnostics: Lower FARs because of the knowledge-based detection technique which produces less false alarms than behaviour-based systems do, which work on models, measurements and thresholds.

The major shortcoming of behaviour-based systems is their high FAR. Especially benign, but yet unknown user behaviour often results in numerous false positives. By that, the number of false alarms can achieve thousands of messages per day in a large network – resulting in an inability to distinguish between true and false alerts.

B. REQUIREMENTS FOR THE ARCHITECTURE

Already a number of general requirements have been derived (e.g. [7, 8]):

- **Security:**
  - *Hidden to the attacker;* in order to analyse the behaviour of an intruder in detail (for instance the exploits used or the steps performed), the presence of intrusion detection tools has to remain completely hidden to the attacker
  - *Multilevel system of interlocking security mechanisms;* despite the careful selection of systems and a consistently focus-oriented security configuration of services, individual security mechanisms can fail (e.g. due to software bugs in the operating system). Through the implementation of additional protective measures at different levels, a multilevel security system still provides protection even if a policy or a device fails.
  - *No influence on productive systems or data streams*
  - *Emergency routines,* e.g. out of band communication to shut down all other communication links
  - *Simulation of a productive behaviour*
  - *Protection of the productive systems*
  - *Recording of all activities*

- **Legal Requirements:**
  - ○ *Prevention of further proliferation of malware and at the same time:*
    - *Minimal and controlled communication* from the research environment to the Internet (especially needed for the analysis of malware behaviour and botnets)
    - *Manipulation of dangerous outgoing data packets*
  - ○ *Consideration of the legal rules and any liability*
- **Scalability:**
  - ○ *Handling X devices, Y users and Z applications*
  - ○ *Evaluation results are independent from data rate*
- **Management challenges:**
  - ○ *Cope with lots of alerts*
  - ○ *Reduce FAR* with the use of intelligent alert correlation [9]
  - ○ *Provide sufficient alert messages* for adequate incident diagnostics
- **Comparison:**
  - ○ Optimization and experimental policies can be tested *without influence on the productive system*
  - ○ *Test behaviour of different IDS based on the same data*

C. EVALUATION CRITERIA

Measuring performance or efficiency of IDS is widely discussed, e.g. Debar et al. give a definition in [10]:

- **Accuracy:** Proper detection of attacks and absence of false alarms. Thus, inaccuracies are anomalies or intrusive traffic flagged as legitimated information.
- **Completeness:** Property to detect all attacks. Without a global knowledge about attacks or abuses of privileges it is much harder to evaluate this measure.
- **Performance:** Rate at which audit events are processed. Real-time detection requires a good system performance.
- **Fault Tolerance:** IDS itself should be resistant to attacks; otherwise new peril points would be opened.

- **Timeliness:** Propagate analysis as quickly as possible in order to react and thus prevent the attacker from harmful malicious actions.

Customers who would like to use the most efficient IDS in their company network have the problem of how to evaluate IDS against the evaluation criteria presented. Therefore, customers have to make a selection of some IDS, implement them, and compare the results afterwards to evaluate the systems concerning their requirements. Several commercial and open-source IDS are available, which make use of knowledge-based or anomaly-based detection methods.

4. ARCHITECTURE

Our proposed architecture divides the network under consideration into several isolated and specialized subnets, namely a Research Network, a Productive Network and an Evaluation Network. At the network border and the border gateway, the datastream is taken of the public network, duplicated one or multiple times and distributed to the different evaluation systems and networks. Test Access Point (TAP) Devices [11], SPAN/Mirrorports and Firewall Kernelmoduls are used for that. Security mechanisms prevent an extravasation from the evaluation networks, e.g., the duplication process is secured by data diodes. By that, only one data direction is possible.

Within the evaluation network, multiple security systems like IDS can be installed, e.g., same systems with different configuration for the optimization of parameters or different systems with complementary detection techniques, e.g. anomaly- and knowledge-based systems. Based on the copied data stream, the evaluation and correlation systems can also initiate active reactions like blocking firewall-ports, generating reports, etc., on dedicated servers therefore not influencing the data on the productive network at all. Also, the results of the different systems can be compared to find the best system respectively configuration for a specific network. Beyond that, the results of different systems can be correlated and, e.g., a majority decision can be taken or more sophisticated correlation techniques can be used or investigated. Figure 2 gives an overview of our architecture. The different components will be described in detail as follows.

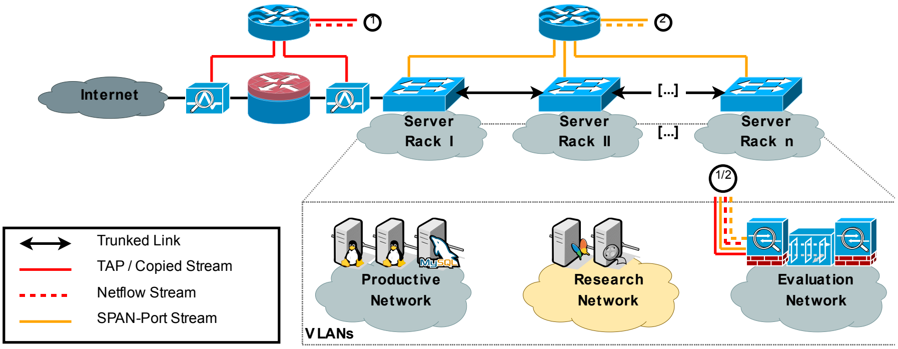

Figure 2. Components of the architecture and integration into company network infrastructure

A. NETWORK BORDER

The border gateway connects all internal networks with the external network, typically an ISP or Internet Backbone. For being able to give basic security against attacks from the external network, a firewall is implemented into the border device. Anyway, because all network traffic is needed for evaluation (not only the filtered one), TAP devices are installed behind as well as in front of the firewall. The TAP devices are able to copy all incoming and outgoing traffic including all Layer 1 and Layer 2 errors. The copied data stream of the TAPs is sent to Cisco Routers which are generating Netflow statistics on the one side; session monitoring is used to multiply the data stream for the evaluation systems on the other side. After the filtering by the firewall, the remaining network traffic is sent through multiple switches which are connected by Trunks/Tagged Ports [12] to the end systems. For complete evaluation and analysis of the internal network, each switch in turn is connected to a second TAP device by its respective monitoring port. The internal network consists of the following parts:

B. RESEARCH NETWORK

This network is used for the examination and evaluation of systems and services relevant to security. To attract attackers and allure malware, honeypots as well as specially prepared real-world servers and services are run in this environment. All systems belonging to the research network can communicate to each other and connections to the Internet are allowed initially. Because of its special nature being open for attacks, only rudimentary or even none filtering is done for this network in the beginning. By an adaption of the rule-set of the firewall, special configurations can be done, e.g., for focusing on individual security aspects.

C. PRODUCTIVE NETWORK

All IT-systems which are required for the operational day-to-day use are set into this network, e.g., office computers, network printers or storage. While the network constraints of this subnet guarantees an usability of the productive network like in the original state without a separation of different networks, multiple security features are activated to protect the corresponding systems:

Switch-based options like fine-grained Access Control Lists (ACLs), Port Security and Community/Isolated port-based VLANs [13,14], which are a basic element for the realization of our architecture. By the use of these elements, combined with a restrictive configuration of the firewall, the productive network can be safely operated in our architecture. Even more, (anonymized) data from the productive network can be mirrored to the research network, therefore generating a more attractive and especially more realistic target environment. Finally, all systems placed in the productive network are monitored by the security systems of the evaluation network, too, enabling an additional protection layer.

D. EVALUATION NETWORK

All network traffic of the complete environment is duplicated for the evaluation network. More precisely, there are even multiple copies of the streams, done by the TAPs in front/after the firewall and the different SPAN-ports of the internal switches. This enables an in-depth evaluation of all network traffic and traffic characteristics, opening up the possibility to build and evaluate complex intrusion and insider detection systems as well as the development of new correlation strategies. The systems of the evaluation network are independent from the systems of other networks, the only receive all data of the other networks but typically without retral information flow because of the TAPs and diodes used. Within the evaluation network, several IDS can be connected or used to exchange evaluation results in order to investigate strategies for reducing the FARs. Based on our architecture, all data as well as corresponding statistics are available for multiple evaluation systems. Even more, the different data streams enable the specialisation of detection systems, e.g., for insider detection, early warning or correlation-based attack detection. Of course, by providing at least the complete external and internal network traffic and the respective flow data, an extensive amount of data has to be processed in the evaluation network. The effective analysis and reasonable storage of parts of the data streams has to be done by the security systems of the evaluation network. Anyway, a central storage of the data is done based on flow data because of the amount of traffic and data protection regulations [15]. Metadata is stored, too; for example, alerts of the different IDS. Based on that, attack sequences can be reconstructed later on.

E. MANAGEMENT NETWORK

For the surveillance of the systems and switches, a separated network is created in our architecture. The systems under surveillance are connected by an independent network interface, which is only used for the management aspects. Also, the systems can be configured and managed by this separated network, e.g., by the use of Nagios and OpenNMS. Attacks based on the management network are prevented based on a strict configuration of the underlying Community VLAN.

F. INTERCONNECTOR

The Interconnector is a coupling device to provide specific communication channels across the different networks. Because of that, the coupling device is particularly critical for security and must be secured especially. To prevent attacks conducted over the coupling device, only minimal functionality is implemented into the device and a comprehensive firewall is integrated. Because of its strict orientation to IT security, the Interconnector is be realized based on OpenBSD [16]. The installation is limited to absolutely necessary packets. For the communication across the networks, OpenSSH is the only service used based on Public Key Authentication.

5. PROOF OF CONCEPT AND CASE STUDY

For the fulfilment of our future research in the area of Intrusion Detection and network security and as a Proof of Concept, a network environment based on the proposed architecture has been realized in our labs. At the moment, the subsequent elements are used for the setup:

- **TAP-Device:** for multiplying the data streams
- **Manageable Switches:** with configured SPAN-Ports for the analysis of the internal network
- **Cisco Routers:** with Netflow capability for the generation of flow data
- **Console Server:** for the management of switches, routers and IPS

Four racks with numerous different servers and systems are integrated in our network at the moment. The connection coming from the Internet is secured by a firewall. OpenBSD is used for this because of its strict orientation to IT security, with a minimum installation of packets and services.

A TAP-device is installed in front of the firewall to be able to capture all in- and outgoing data. For each connection conveyed through the TAP, multiple copies of the data stream can be tapped with separated channels for the incoming and outgoing network packets. Next, the data stream of the TAP is sent to the evaluation network and to Cisco Routers with Netflow Capability for the generation of respective statistical data which is also sent to the evaluation network. Also, additional copies of the data stream can be generated via monitoring. The regular data stream after the firewall is routed by manageable switches which are trunked together and configured for the different VLANs in use. The configuration of the firewall respects the structure of the VLANs, e.g., the traffic to the research network typically is not filtered while traffic to systems of the productive network is altered and controlled. To enable an in-depth view onto the internal network, for example for the detection of Insider Activities, each manageable switch is configured to sent all traffic to a SPAN-Port.

All SPAN-Ports are collected by the evaluation network and used for Insider and Extrusion Detection. Because the amount of data transported in the different segments of the internal network is not as high as the traffic volume running through the external interfaces to the Internet, the SPAN-Ports are typically able to copy all data without a loss of packages. Note, that this is not possible when all ports of a switch are heavily used: Because the SPAN-Ports are “regular” ports with the same capacity like all other switch-ports, packets will be dropped randomly in case the volume of the aggregated packets is higher than the data rate of the SPAN-Port. Therefore, this technique can be used to reliably supervise individual network segments, e.g., the systems of one rack, but can’t be used at the borders of the network where TAP devices are needed. It also should be mentioned, that trying to increase the traffic on one switch to force a high amount of dropped packets (e.g., to conceal an attack) is not possible, because this results in a strong deviation to the normal network behaviour, and thus easily detectable by behaviour-based IDS. The data stream of the SPAN-ports is sent to Cisco Routers for the further distribution and the generation of Netflow data.

At the moment, the analysis and evaluation of the Intrusion Detection is done with the following systems (amongst others) in particular:

- **Cisco IPS 4345** (Signature Release S690, January, 16th 2013)
- **Snort Version 2.9.3.1** (snapshot-ruleset downloaded by PulledPork / SourceFire VRT rules, December, 20th 2012 and Emerging Threats rules)
- **Bro Version 2.1**, including snort2bro-translated signatures
- **Suricata 1.3**, Emerging Threats and SourceFire VRT rules
- **Flowmatrix Version 0.30**

Because most systems are running on Virtual Machines in the evaluation network, additional systems can be integrated fast and easily. Furthermore, the data stream can be anonymized and saved into a database in the evaluation network for, e.g., repeating experiments or optimizing IDS-parameters.

One of the advantages of our environment is the possibility of testing and optimizing new systems and configurations without an endangerment respectively disruption of the productive network. For example, when installing the Cisco IPS 4345 device, all rules for blocking traffic had been disabled. Even so, after putting the system into a TAP-link, it started to drop http-traffic. Incidents like this can have serious consequences when they interrupt systems and services in productive networks. In contrast, deployed inside the evaluation network and only working on the copied data stream, unwanted effects cannot influence the productive network and an evaluation of systems and their configurations is possible.

Another aspect is the system performance of IDS in real-world environments. As studies, e.g., by NSS-Labs have shown, often systems are not able to fulfil the specified performance. Until now, our evaluations produce similar results. Several of the considered systems have produced multiple unreported errors during runtime, not recognizable with their User Interfaces. In some situations, systems dropped up to 95% of the network packets or ended in a very high, incomprehensible use of system resources. We will investigate these phenomena in more detail as part of our current research because it is a crucial factor for the utility of IDS in real-world networks.

Even when the systems are running as expected, numerous False Alerts are hampering the use in today's networks. Therefore, our architecture can be used to operate multiple IDS in parallel without any interference or endangerment of the networks. The variety of systems can be used to develop and evaluate correlation strategies aiming for an improvement of detection- and false alarm rates.

At the moment, we are running multiple IDS for the evaluation of security incidents. By that, we have systems specialized and configured for four different doctrines:

- **External Attack Detection** on the Border Gateway
- **Insider Detection** on the Internal Network
- **Misconfiguration Detection** on all Networks
- **Data Leakage Detection** on all Networks

New algorithms and techniques for Alert Correlation are currently under our development. Figure 3 gives an overview of the system used for this purpose.

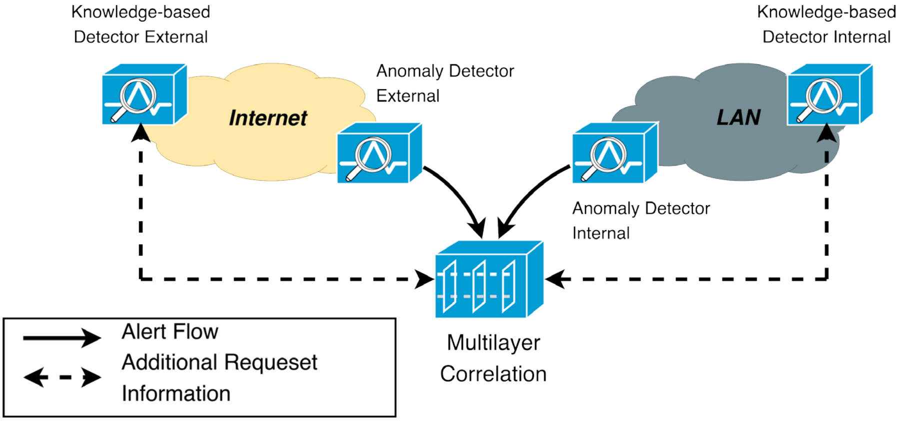
Knowledge-based Detector External
Internet
Anomaly Detector External
Knowledge-based Detector Internal
LAN
Anomaly Detector Internal
Multilayer Correlation
Alert Flow
Additional Requeset Information

Figure 3. Multilayer-correlation of Sensor Information

The correlation between behaviour-based respectively knowledge-based sensors on different layers (e.g., internal and external network) on the one side and the correlation between behaviour-based and knowledge-based sensors on the same layers is of special importance for improving the detection accuracy and lowering FARs. Therefore, a central component placed in the evaluation network collects alert and sensor information from the different IDS of our network. Based on the doctrine under evaluation, alert and sensor information of different components is correlated and afterwards, further information is requested from other systems. For example, for the External Attack Detection, the behaviour-based alerts of the sensors in front of the border gateway and after the gateway inside the internal network are correlated. By that, chances are increased that a new, unknown attack, which cannot be detected by knowledge-based systems, can be found by the correlation of external and internal deviations and FARs can be reduced. On the other side, further information can be generated by the use of other sensor information, e.g., to narrow down external alerts. See Figures 4 and 5 for an example.

As illustrated, the typical situation, a higher number of anomalies in the external network and a lower number of anomalies in the internal is the case. Often, these anomalies are based on benign, but yet unknown user activities, therefore generating False Alerts. By the correlation of external and internal alerts, events of high relevance can be identified and checked.

For example, the calculation of standard deviations of characteristic traffic parameters can be used to manually identify low intensity anomalies. The IDS Flowmatrix gives a graphical representation, which enables the operator to visually identify anomalies (e.g., special patterns of higher deviations), which are below the regular alert thresholds of the IDS.

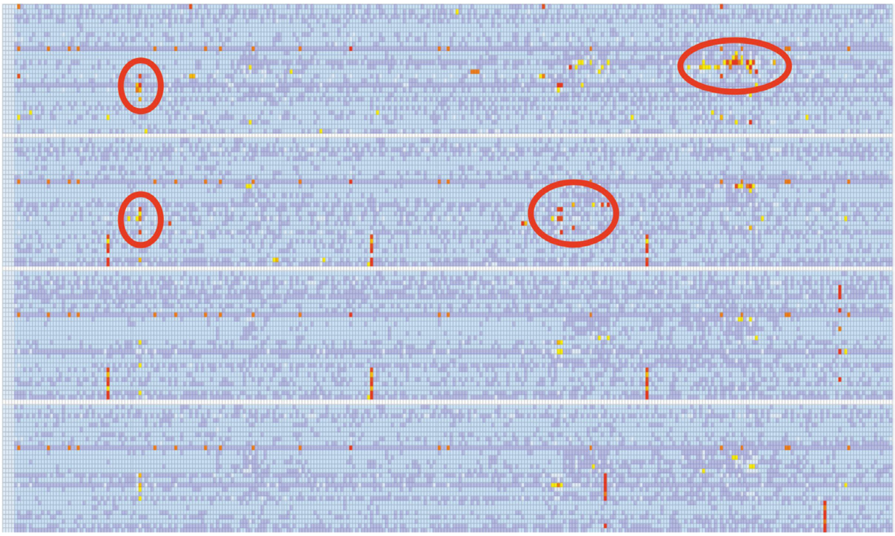

Figure 4. Degree of Standard Deviation for IP Addresses and Ports in the External Network. Clusters of interest are marked

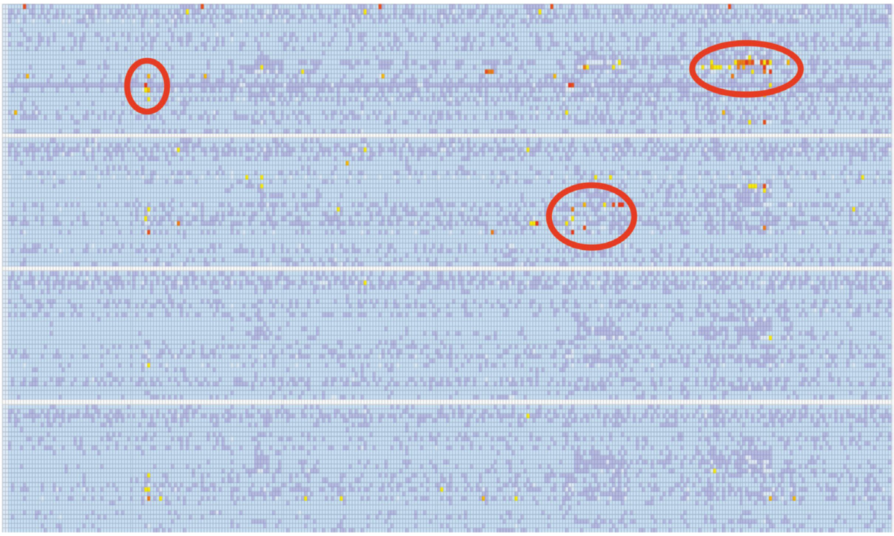

Figure 5. Degree of Standard Deviation for IP Addresses and Ports in the Internal Network. Clusters of interest are marked

Figures 4 and 5 present two screenshots of the analysis of a 300-minute timeslot done by Flowmatrix [17], for the external as well as the internal traffic. The warmer the colour, the higher the deviation of a parameter within a cluster; which is a sign for an anomaly in the corresponding dataset. The next step of our ongoing work will be the development of correlation techniques, which are able to use such kind of information to reduce FARs on the one side and detect sophisticated attacks on the other.

Based on this knowledge about an event of interest, further data can be collected from other sensors. For example, based on the observation time and IP addresses, collected flow- and header-information can be analysed or knowledge-based systems can be checked for suspicious log-entries. If an identification of an attack is possible, the collected information can be used for a rapid and (semi-) automatic development of new patterns, which is a further goal of our research.

It is important to differentiate if and which alerts are seen in front and behind the border gateway: Typically, external and internal alerts will arise in case of a successful External Attack. Here, more events will be registered on the external network, but often with low intensity, e.g., a service scan. On the other side, after breaking into the system, the attacker will try to investigate the internal network inconspicuous. Because deviations are more significant in the controlled internal network, it is easier to filter our events of interest. These information can be used again for the selection of the relevant events in the external network, which otherwise don't exceed thresholds or decline in the background noise.

The need for intelligent new correlation techniques can be seen in Figure 6.

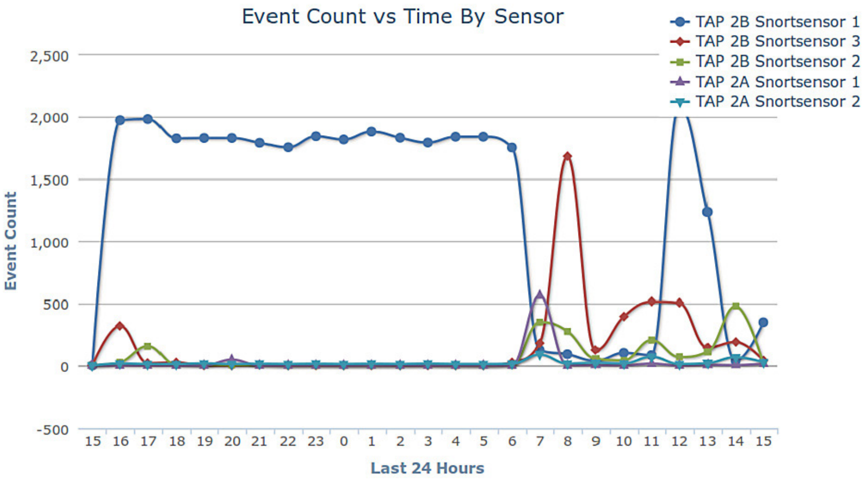

Figure 6. Events registered over 24 hours by different Snort-Sensor in an Academic Network

As depicted, the events produced by the different sensors are quite irregular. Even throughout the night, a lot of alerts had been generated. Such a strongly irregular behaviour is difficult to learn by a behaviour-based system when the underlying model has to be created. Often, these alerts are harmless and caused by inoffensive actions. Therefore, it is necessary to reduce these alerts by our proposed multilayer-correlation. Based on that, this knowledge also can be used to optimize and control the learning phase of behaviour-based IDS, for example by the pre-filtering of unwanted traffic (which is otherwise learned as normal behaviour and therefore cannot be detected later on in the operational mode).

6. CONCLUSION

Network analysis in company networks is used frequently to discover potential attacks on the computer network. Traditionally, especially IDS are used to defeat these attacks. But often information system departments have difficulties to decide which of the offered IDS they should use. It is often an open question if the system will really fulfil its individual requirements such as different network structure, offered and used services, etc. and therefore companies tend to setup a test implementation before buying the product. One of the biggest problems is that a fair comparison (ceteris paribus) is not possible unless the test environment is equal. Today the well-known DARPA data-set is still used to compare IDS [3]. Because of multiple design errors, the data-set was often criticized and scientists disadvise using it any longer [18, 19]. Even if there are other data-sets available (e.g. MAWI Working Group [20], GEANT [21], ACM SIGCOMM [22]), up to now, none of them was able to prevail. Even more, it has been shown that there is often a strong difference between synthetic- and real-world data based evaluations of IDS (e.g., see [8]). Instead of using fixed and outdated data-sets, our concept shows the possibility to compare systems based on real world data. As described in the proof-of-concept section, our proposed architecture is inserted transparently into the productive network. Thus the architecture gives the opportunity to capture real-world data in real-time as well as the possibility to provide different IDS environments with the same data set for a fair evaluation as well as a sophisticated correlation.

Our architecture is transparent and allows several IDS environments to be implemented in parallel which can be used for configuration optimization, error checking, monitoring the learning phase of anomaly based IDS, etc.

The architecture has also withstood attacks when different security vulnerabilities became known (and exploitable), such as “Multiple Vulnerabilities in Cisco Firewall Services Module” [23] concerning the Cisco firewall module of one of the core switches used. Due to the multi-layered security, attacks in this case were already effectively blocked both by the ACLs of the access switches and the firewall.

The next step will be an integrated and common graphical user interface for the configuration, selection and surveillance of the IDS as well as the fast and easy specification of the rulesets used for the correlation of IDS events. Based on that, the advantages of different system architectures and capabilities can be combined and synergetic effects can be enabled, generating more reliable and secure IDS. Therefore, we are planning to include communication standards and mechanisms like, for example, the data exchange by the Common Intrusion Detection Framework (CIDF) [24], the Intrusion Detection Message Exchange Format (IDMEF) and its associated protocols (Intrusion Detection eXchange Protocol (IDXP), Intrusion Alert Protocol (IAP), Blocks Extensible Exchange Protocol (BEEP); [25, 26, 27]) or the Intruder Detection and Isolation Protocol (IDIP) [28, 29].

Even if a system is not able to provide one of these standard mechanisms, log and alert-files can be evaluated by the use of regular expressions in an efficient way, opening the possibility for integration as well.

Another important aspect of our future work is the conception and development of a correlation strategy for the integrated IDS. As already shown, numerous aspects must be taken into consideration when correlating alerts: For example, as shown by the evaluation of the prototype, some (correct) alarms are only raised by single systems, therefore a majority decision is not enough. Our test environment provides the basis for the development of required, sophisticated correlation techniques.

REFERENCES

[1] Winterfeld, S., & Rosenthal, R.: Understanding Today's Cyber Challenges. Manager, (703), 20151. Retrieved from http://www.tasc.com/news_media/white_papers/TASC_Cyber_Challenges_Study_May_2011_FINAL.pdf, 2011

[2] NSS Labs: Network IPS Group Test 2010. https://www.nsslabs.com/reports/network-ips-group-test-2010, 2010

[3] Lippmann, R., Haines, J.W., Fried, D.J., Korba, J., Das, K.: The 1999 DARPA Offline Intrusion Detection Evaluation. Tech. rep., Lincoln Laboratory MIT, 244 Wood Street, Lexington, MA, 2000

[4] Cuppens, F., Autrel, F., Miege, A., Benferhat, S.: Correlation in an intrusion detection process. Proceedings SEcurite des communications sur internet (SECI02) pp. 153-171, http://www.lsv.ens-cachan.fr/~goubault/SECI-02/Final/actes.pdf\#page=153, 2002

[5] Debar, H., Wespi, A.: Aggregation and Correlation of Intrusion-Detection. In: Recent Advances in Intrusion Detection, Springer 2001

[6] Valeur, F., Vigna, G., Kruegel, C., Kemmerer, R.A.: Comprehensive approach to intrusion detection alert correlation. In: Dependable and Secure Computing, IEEE Transactions on. . pp. 146-169. IEEE 2004

[7] Golling, M., Stelte, B.: Requirements for a Future EWS – Cyber Defence in the Internet of the Future, 2011 3rd International Conference on Cyber Conflict, IEEE, 2011

[8] Koch, R.: Towards Next-Generation Intrusion Detection, 2011 3rd International Conference on Cyber Conflict, IEEE, 2011

[9] Boggs, N., Hiremagalore, S., Stavrou, A., Stolfo, S.: Experimental Results of Cross-Site Exchange of Web Content Anomaly Detector Alerts, IEEE Conference on Technologies for Homeland Security, Boston, 2010

[10] Debar, H., Dacier, M., Wespi, A.: A revised taxonomy for intrusion-detection systems, Annals of Telecommunications 55(7), 361-378, 2000

[11] Network Working Group and others: IETF Policy on Wiretapping. RFC 2804, May 2000

[12] IEEE: 802.1 Q/D10, IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks, Copyright by the Institute of Electrical and Electronics Engineers. Draft, 1997

[13] Hucaby, D., McQuerry, S.: Cisco Field Manual: Catalyst Switch Configuration. Cisco Systems, 2003

[14] Cisco Systems. Inc: Configuring Isolated Private VLANs on Catalyst Switches. Cisco Systems http://www.cisco.com/image/gif/paws/40781/194.pdf, 2008

[15] Claise, B.: RFC 3954: Cisco Systems NetFlow Services Export Version 9. IETF http://www. ietf. org/rfc/rfc3954. txt, 2004

[16] Cowan, C.: Software security for open-source systems. In: Security and Privacy, IEEE Transactions on. . pp. 38-45. IEEE 2003

[17] Xharru Ltd. AKMA Labs: FlowMatrix - Network Behavior Analysis System, http://akmalabs.com/flowmatrix.php, 2012

[18] McHugh, J.: Testing intrusion detection systems: a critique of the 1998 and 1999 darpa intrusion detection system evaluations as performed by lincoln laboratory. ACM Trans. Inf. Syst. Secur. 3(4), 262-294, dOI http://doi.acm.org/10.1145/382912.382923, November 2000

[19] Athanasiades, N. et al.: Intrusion detection testing and benchmarking methodologies. In: Information Assurance, 2003. IWIAS 2003. Proceedings. First IEEE International Workshop on . IEEE Computer Society, 63-72, March 2003

[20] MAWI Working Group Traffic Archive. Website, http://mawi.wide.ad.jp/mawi/, last seen on February 24th, 2012

[21] MoMe Cluster of European Projects aimed at Monitoring and Measurement. MOME Database. Website, http://www.ist-mome.org/database/, last seen on February 24th, 2012

[22] ACM SIGCOMM. Internet Traffic Archive. Website, http://www.sigcomm.org/ ITA/, last seen on February 24th, 2012

[23] Cisco Security Advisory: Multiple Vulnerabilities in Cisco Firewall Services Module, Cisco Systems, http://tools.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-20121010-fwsm/, 2012

[24] Kahn, C., Porras, P., Staniford-Chen, S., Tung, B.: A common intrusion detection framework. Submitted to Journal of Computer Security, 1998

[25] Rose, M.: The Blocks Extensible Exchange Protocol core, 1–59. Retrieved from http://www.hjp.at/doc/rfc/rfc3080.html, 2001

[26] Corner, D.: IDMEF-Lingua Franca for Security Incident Management Tutorial and Review of Standards Development. SANS Institute, 2003

[27] Rose, M.: Mapping the BEEP Core onto TCP, 1–9. Retrieved from http://tools.ietf.org/html/3081, 2001

[28] Schnackenberg, D., Djahandari, K., Sterne, D.: Infrastructure for intrusion detection and response. In: DARPA Information Survivability Conference and Exposition, 2000. DISCEX'00. Proceedings. vol. 2, pp. 3-11, IEEE 2000

[29] Kothari, P.: Intrusion Detection Interoperability and Standardization. SANS Institute. Retrieved from http://cs.uccs.edu/~chow/pub/master/sjelinek/doc/research/idmef.pdf, 2002