



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Data Mining: A New Intrusion Detection Approach

GIAC Security Essentials Certification Practical Assignment

Version No 1.4 Option 1

By: Dary Alexandra Pena Maldonado

June 19th 2003

Introduction

The well-known use of Intrusion Detection Systems (IDS) has revealed the significance of their key role in network security as well as its evident performance shortfalls. For the last decade, vendors and researches have found on Data Mining (DM) a helpful practice to uncover new insights, associations and hidden patterns within large data set of logs and messages. Starting from the current IDS state of art, this paper outlines how DM techniques have been applied conveniently in this context including the required architecture, some of the obtained results; as well as some directions about its future application.

IDS State of Art –An Overview-

The size and complexity in enterprise networks, exposed mainly by their geographical distribution, the diversity of technologies, linked partners (trusted and non-trusted) and security policies, make them challenging from the Information Security point of view. In fact, these features are associated directly to the scalability of their security architecture (responsible for the protection of a large number specific endpoints: users, hosts, servers, domains, network segments); and the obstacles revealed during its setup, management and troubleshooting.

Intrusion Detection Systems (IDS) have become a crucial element to secure their current and emerging networks, as well as their services and applications detecting, alerting and responding to malicious activity. Organization can use IDS employing as **Data Source** the information from an individual host (*Host Intrusion Detection HID*) and/or network of computers (*Network Intrusion Detection*). On the other hand, it could be selected according to their **Detection Strategy** finding intrusions that matches with pre-defined patterns or signatures (*Misuse Detection*) and/or finding intrusions by the expected network behavior and its deviations (*Anomaly Detection*)

Efficient deployment of IDS in Global Organizations is a real challenge especially with the selection and implementation of any ID technique and product(s) that fits their requirements according to their *technological environment* (Security policies, network architectures, security management capability) and the *state of practice of commercial IDS*. This condition and the current security lacks reveal some issues that have to be considered regarding to the IDS development and their use:

- The growing index of reported incidents¹, their complexity and their changing nature has been increased notoriously making unsuitable the exclusive use of conventional ID Systems. Additionally, the advances in the automatic tools and strategies needed to perform malicious of activity, enable to non-technical individuals in the process of automating more complex attacks.
- The spectrum of employed IDS techniques and products (from multiple vendors), in addition to the lack of an accepted universal standard concerning to their integration makes things worse when the security analysts have to effectively analyze, monitor and *correlate* this data coming from distributed closed and proprietary components
- Considering that an Intrusion Detection System is basically a flow and packet inspector, some issues could prevent them of seeing all traffic needed for the detection:
 - The bandwidth of both internal and external enterprise networks is increasing notoriously (typically from fast Ethernet to Gigabit Ethernet), making as a requirement the improvement on their current computing performance. To succeed these traffic loads, it needs to reach acceptable packet-processing and per-packet latency rates². Based on this principle, an intruder could exhausts resources as CPU cycles, memory and/or network bandwidth forcing the IDS to process non-useful information and dropping packets.
 - The usage trend of switched networks, modems, VPN's and encryption technologies facilitate the malicious traffic transport inside the network masking the packets to the IDS
- Marketable ID Systems frequently use string matching patterns (with/without small variants) detecting known intrusion patterns, while gather information during all the attack time line³. However they have limited correlation capabilities between pre-attack and attack diminishing the analyst reaction capability just before that the thread of intrusion becomes damage. In fact, they are more focused in detection more than in prevention.
- After the damage, commercial IDS don't provide enough specific information about the issue scope, its possible solution and the ability to get valid data legally to evidence the attacker action. A large number of IDS record processed events more than network packets.
- Common ID System architectures consist in a central management unit that receives information from sensors that collect activity individually; focusing on low level attacks and alarming independently without

considering logical connections between them. They need to maintain a complete view of the network, correlating traffic and alarms from different sources including NIDS, HIDS, firewall, router logs, etc.

- Enterprise environments can evaluate their ID Systems with a number of **Sent Attacks**⁴ and associated merit figures as **True Detection Rate** based on: **True positives** (Successful attack detecting index), **False Positives** (detecting and signaling index of attacks occurred when there is no attack) and **False Negatives** (failing index to detect and signal an attack that has occurred or inability to detect new types of self-modifying attacks). Considering that the amount of their traffic makes more complex and extremely time consuming the analysis performed by the technical staff, choosing and deploying an adequate IDS solution enables the generation of manageable number of false positives⁵ and false negatives. It will support credible conclusions or diagnosis provided and ensuring the effectiveness to the intrusion responses.
- Management activity over ID Systems requires a continuous and extensive dependence of high skilled analysts, their perception, reasoning and experience, with the associated high risk and cost of ownership (deployment and maintenance).

Although the usage trend of Intrusion Detection Systems has been widespread as a response to the increasing number of incidents by System Penetration [Figure 1], these types of performance issues are limiting the IDS scope to small and non-cost sensitive deployments weakening their add value for the defense against the increasing malicious activity. The requirements overcome the scope of this IDS generation, facing to the industry and the researches to find more refined and efficient detection methods. Migrate from Intrusion Detection to Intrusion Prevention will ensure cost savings in security monitoring and management.⁶ This set of issues has given to Data Mining the chance to contribute as a new approach for the Intrusion Detection field.

ID Systems vs Attacks by System Penetration

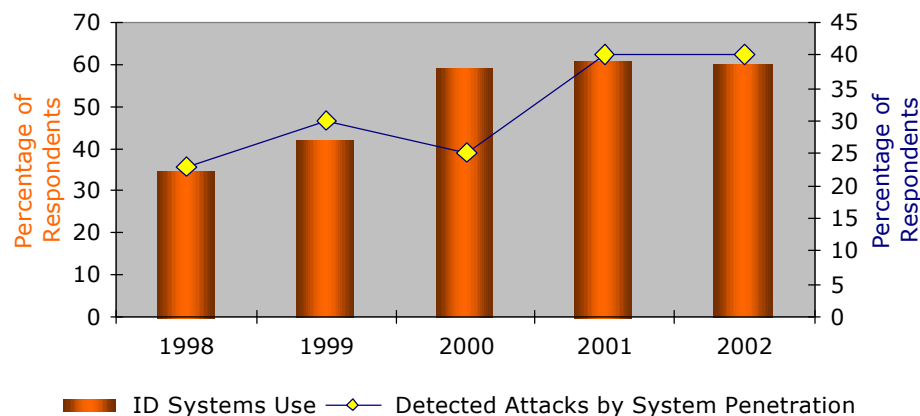


Figure 1. ID System use and Detected Attacks by System Penetration Trends⁷
Knowledge Discovery And Data Mining Outline

Technically the *Knowledge Discovery in Databases (KDD)* practice is associated with the extraction and discovery of useful information from large relational databases while *Data Mining (DM)* represents its core as decision support stage [Table 1]: Data Mining is a finding process of significant non-intuitive correlations and patterns from a variety of sources, making possible to get high-level knowledge information from low-level data.




Figure 2. Discovery Knowledge Process⁸

Stage	Comments
Data Selection and Cleaning (Preprocessing)	Selection of relevant items for representation <ul style="list-style-type: none"> •Integrate data from different formats and Heterogeneous sources •Verify correct ranges and limits Data Cleaning of invalid data <ul style="list-style-type: none"> •Use existing data to extrapolate •Compute the least noisy representation
Transformation (Preprocessing)	Input-format related transformations
Data Mining / Pattern Discovery	<ul style="list-style-type: none"> • Identify sets of information to be mined. • Technique Selection: association mining, Clustering, time-series analysis. • Extend the patterns to larger selections • Pattern Evaluation
Interpretation Knowledge(Post Processing)	<ul style="list-style-type: none"> •Documentation •Order patterns with metrics. •Highlight most relevant patterns •Visualization

Table 1. Discovery Knowledge Process

Their applications in Large Organizations have been successfully widespread as helpful and adaptive tools to analyze Survey Results, Fraud Detection,

Manufacturing Process, Risk Analysis and Management, Market, Sales, Scientific Data and network management studies relating its scope with the appliance of related disciplines as machine learning Database Systems, Data Warehouses and Statistics.

Data Mining And Intrusion Detection Systems

Recently, researchers and vendors have explored Data Mining as a new valuable approach to the IDS. They are exploring the chance to improve data reduction, discovery and detection capabilities covering hidden patterns, deviations of known attacks or new ones; and maximizing the cost-benefit relationship for an ID deployment.

ID using Data Mining (IDDM)⁹, use as basis the audited data from different sources (particularly records representing a network event, described with attributes as number of bytes transferred, access counts, etc), activity indexes (from normal and intrusion activity) and algorithms to search significant patterns; enabling the construction of misuse and anomaly detection models based on an intelligible set of rules. The raw data [Figure 3] is archived and sampled in discrete records according to the attributes. Data mining programs are subsequently used over the traffic records to compute patterns. The connections and the patterns are then analyzed to construct additional features, getting an empirical and iterative approach. One of the most critical and success determining selections is the related with the data mining technique:

- **Classification** categorizes the data records (training data set) in a predetermined set of classes (Data Classes) used as attribute to label each record; distinguishing elements belonging to the normal or abnormal class (a specific kind of intrusion), using decision trees or rules. This technique has been popular to detect individual attacks but has to be applied with complementary fine-tuning techniques to reduce its demonstrated high false positives rate.

With support tools as RIPER (a classification rule learning program) and using a preliminary set of intrusion features, accurate rules and temporal statistical indexes can be generated to recognize anomalous activity. They have to be inspected, edited and included in the desired model (frequently misuse models).¹⁰

- **Association Rules:** Associations of system features finding unseen and / or unexpected attribute correlations within data records of a data set, as a basis for behavior profiles.
- **Frequent Episode Rules** analyze relationships in the data stream to find recurrent and sequential patterns of simultaneous events, to compute them

later. Its results have been useful for attacks with arbitrary patterns of noise or distributed attacks.

- **Clustering** discovers complex intrusions occurred over extended periods of time and different spaces, correlating independent network events. The sets of data belonging to the cluster (attack or normal activity profile) are modeled according to pre-defined metrics and their common features. It is especially efficient to detect hybrids of attack in the cluster, showing high performance when are processed features computationally expensive. With other techniques is able to re-train itself reclassifying the existing clusters and generating new ones.
- **Meta-rules** derive change rules over a period of time, comparing the status of two data sets and describing their “evolution” in time based on their common, modified and new features.




Figure 3. ID Process using Data Mining¹¹

Typically, the best features for these techniques and detection models are combined to obtain a high detection performance and a complex profile for intruders. Recognized practices merge models for new activity (attacks or normal events) and the existing models, to generate adaptive processes able to learn inductively the existing correlations: this Meta-Learning capability and its adaptability with other techniques, as classification and association rules, have been evaluated empirically as effective and scalable. The rules reduce substantially the impractical manual development process of patterns and profiles, computing statistical patterns from the collected data.¹²

Data Mining can be applied to misuse detection and anomaly detection models. In **Data Mining-based Misuse Detection** each data record is classified and labeled as normal or anomalous activity. This process is the basis for a learning algorithm able to detect known attacks and new ones if they are cataloged appropriately

under a statistical process. The basic known as discovery outliers, matches an abnormal behaviour against an attack patterns knowledge base that capture behavioral patterns of intrusion and typical activity. To do this, it is needed to compute each measure with random variables implying more updating effort as more audit records are analyzed but more accuracy with more mined data. Although the activity needs to be analyzed individually, complementary visualization and data mining techniques can be used to improve performance and reduce the computational requirements. Some researches focused on this topic are JAM (Java Agents for Metalearning), MADAM ID (Mining Audit Data for Automated Models for Intrusion Detection) and Automated Discovery and Concise Predictive Rules for Intrusion Detection¹³.

On the other hand, the **Data Mining-based anomaly detection** goals are related with searching inherent but previously unidentified information from the collected data. A set of records is stored building a normal profile to be compared with the most recent activity (delimited in a time window) determining if it is far from the expected behaviour and establishing the similarity degree with other historical profiles. The suspicious connection is classified as known, unknown or false alarm. The most popular anomaly detection system using data mining is ADAM (Audit Data Analysis and Meaning).

Some projects have experimented with a hybrid approach that uses an artificial anomaly generation method, allowing supervised inductive methods for anomaly and misuse detection providing synthetic alarms.¹⁴

Architecture Requirements

Data Mining is itself a consuming computing resource process. An efficient deployment for an ID System employing Data Mining requires an adaptive and scalable *architecture* and *infrastructure* able to support the storage for the audited data, its processing, the model generation and distribution, as well as the interaction with the pre-existing elements in the organization security infrastructure.

The [Figure 4] shows a modular approach to this architecture using a model proposed by Tomas Abraham¹⁵. According to specific applications and environments, particular modules or elements (i.e. learning and detection agents) should be added in their context, complementing the task of the selected data mining technique.

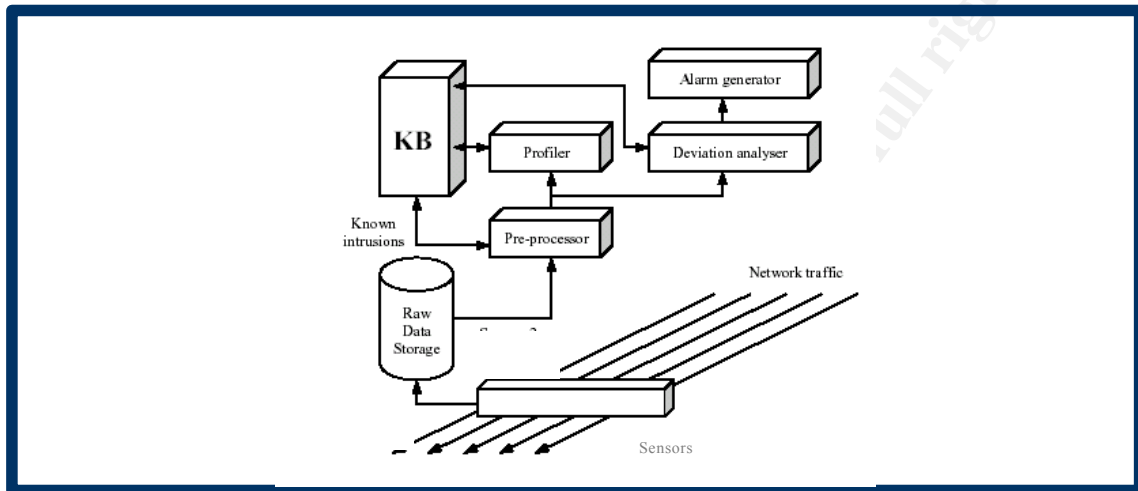
The Data Mining operation to detect fresh attacks or avoid additional damage in an ongoing attack requires hardening of the current infrastructure including atypical elements additional to the needed in an ordinary Intrusion Detection

System deployment:

- Sensors with optimal performance regarding to its data volume management and acquisition rate. Packet sniffers as NFR or Bro have been tested in researches projects.¹⁶
- A high performance production database efficiently designed, administrated and tuned, responsible to maintain the data and model updating process as well as retrieve efficient query results and record alarms simultaneously with an acceptable answer time and extensive indexing capabilities. Depending of the database size, the complexity in the queries and the required performance could be needed implementing Massively Parallel Processors (MPP).
- Enough store space to handle the data associated to the data mining process. (Intermediate and final results, temporal files, etc).
- High Computing performance able to support the CPU and memory requirements for data mining process. Some practices recommend at least four times the resources needed to implement a non-data mining-based IDS.¹⁷
- Data Mining specialized software and accessory tools to analyze volume of information enabling to the analyst to extract higher level knowledge from filtered and formatted data
- The communication channels between the sensors and the central archive would be capable to handle large volumes of data guarantying secure traffic.
- The transactions have to guarantee the data security and a redundant scheme would be desirable

On the other hand and considering that human intervention is required in the process, is desirable that the involved staff has the ability to face issues for security, databases and data mining applications.

New researches use meta-learning and third party detection tools to balance the finding patterns process, enlarging the capabilities of data mining as off-line process near to a real on-line process. In actual fact, they are not strictly needed to deploy the system, but they would allow that data could be collected and the rules generated, after a stabilization period of time, in an adaptive and flexible process maintained by the analyst with periodic inspections to the rule set and updates to the Knowledge Base.



Component	Short Description
Sensors	<ul style="list-style-type: none"> •Collect raw network data
Raw Data Storage	<ul style="list-style-type: none"> •Archives the collected in a relational database data and apply filters if it's needed i.e. Transaction Filters by source/destination IP/ port/ protocol with hits in a specific time window •Support creation and tracking for security incidents
Pre-processor	<ul style="list-style-type: none"> •Transforms data in useful formats for mining algorithms. •Allows the use of programmable and customized models. •Filters and eliminates noise. •Uses detection models to recognize known attack patterns storing them in the KB for more analysis and report.
Knowledge Base	<ul style="list-style-type: none"> •Stores mining models, rules and associated information. The data can be manipulated for offline analysis, training and labeling. •Makes easy the data correlation from a variety of sources or collections from longer periods of time, enabling the detection of large scale attacks
Profiler	<ul style="list-style-type: none"> •Get status of set of data in a specific period of time for deviation analysis. •Uses historical information and the data set to generate new profiles, constructs features and redistributes the profiles
Deviation Analyzer	<ul style="list-style-type: none"> •Finds differences to be analyzed storing them in the knowledge base to obtain new profiles computed from the models and triggers the alarms.
Alarm Generator	<ul style="list-style-type: none"> •Notify to the administrator abnormal network events using e-mails, alarms

Figure 4. A modular architecture for Data Mining-based ID System

Most of the importance of this architecture resides on the implicit complexity associated with the amount of data that has to be processed (millions of records per week) and the features provided to get an efficient training of the system as basic setup. Additionally other features as the profiling process, the group of DM techniques selected, the integration with typical IDS products, and a distributed environment could add difficulty in the architecture proposed.

© SANS Institute 2000 - 2005, Author retains full rights.

Success Deployments –Alarm Management-

One of the most interesting contributions of Data Mining in the Intrusion Detection field is the support to the alarms investigation, specifically applied to large organizations and security providers offering security-outsourcing solutions. The context around their activity (many sensors installed in a distributed and different networking environment, heterogeneous traffic, high rate of false alarms, significant manual intervention, the complexity for an appropriate tuning process) and the shortfalls for the current detection scheme, have found on Data Mining-based IDS an option to improve the organizational IDS management. It mines historical alarms and compresses huge logs into highly useful abstracts that enable a more efficient future alarms use.

Most of the success of this process is based on a periodical “cleaning” practice in which the root causes for events triggering alarms are identified, classified (in clusters or groups), and eliminated or at least reduced. In fact almost 90% of the false positives are originated for a small number of root causes and typically are related with configuration issues; so this set of actions removes the associated redundancy and the alarm load linked to non-malicious activity. When is not possible to eliminate a root cause (considering its operative cost or action domain), filtering or correlation rules can be used under the associated cost by loss information, growing in a further insight knowledge.

Exploiting the monotonous alarms nature, some of the most promising results have been achieved using *clustering* combined with *frequent episode rules*, in an alert merging function. Clustering joins all the alarms (from the same or different IDS) as result of the same attack occurrence or the attacks associated to a specific alert, using the similarity concept as integrator element. When a new alarm is triggered the function determines all the related alerts, to create a more general and representative alert for the whole cluster and then apply correlation between the individual elements discovering possible plans for a distributed and complex attack. Focusing in this knowledge, the targets exploited until now are reported and using extrapolation anticipate future intruder intents taking the proper administrative actions. It could be considered as a first step to get intrusion prevention more than intrusion detection.

As a supplementary practice after some time the results become shorter, significant, intuitive and highly interpretable reducing the analysis time and the overall cost per month. Other key alarm characteristic is the mixture of attributes in their messages (categorical, numerical, time, string). It points to an inherent complexity derived from a multi-dimensional space alarm model and the similarity concept used in the cluster building process.

These efforts to improve the performance over the existing IDS's have showed

potential results regarding to the high alarm load reduction (almost 75% of alarms processed automatically). Inherently some additional benefits about this correlation process include a significant getting cost effectiveness, as well as better capabilities to analyze post-mortem events.




Figure 5. Alarm load reduction using a data mining approach¹⁸

In this way the process is directed toward to a full-automated environment close to real-time (collect, correlate and analyze the alarms before a damage occurs), where a knowledge-based system could be included to map a group of alarms to their root causes and an extensive prior knowledge is not needed at the user side; even more when the data-mining role can be complemented with interesting visualization tools.

Specifically IBM has led this kind of deployments for a variety of clients as part of its Internet Emergency Response Service (IERS). IBM integrates conventional third party sensor applications (Symantec / AXENT, Cisco, ISS, Dragon, Sanctum AppShield and Tripwire¹⁹) and offers real-time intrusion detection (RTID), using Data Mining and event correlation to manage the huge alarm volume generated for all the customers. Their results have showed a 90% alarm reduction over more than 100 million of them. This practice allows acting proactively using early attacks detection, as well as the use as forensic and fraud detection tool. After some evaluations, the Incident Emerging Response has resulted in a Great, flexible customer and low maintenance; with a high associated cost (Three years cost equivalent to hiring a security guru) and some overhead problems with large-scale attacks²⁰

Some Commercial & Related Products

Most of the commercial tools and products available for data mining are commonly addressed to marketing and financial markets. Regarding to IDDM, some of the most published products as Intrusion Secure Net²¹ and Tivoli Decision Support²² have used data mining capabilities with a special interest to improve their reporting and forensic potential more than their a real-time detection enhancement: The records are inserted using an internal event cache in a relational database for long-term storage and analysis before the real-time processing. It enhances its aggregation capabilities taking advantage from some data mine features included in commercial database engines as SQL and Oracle.

In a related effort, other researches²³ and vendors like GuardedNet, intellitactics, net Forensics, Open Systems Private I²⁴ and Internet Security Systems (ISS)²⁵ are exploring the chance of use pure correlation to get similar objectives as the purposed by IDDM: beginning from log central storage with predetermined query possibilities until refined reporting capabilities taking different sources based on rule-based and/or statistical correlation; and resulting in a more comprehensive log management. Although the scope can be comparable, the conceptual basis is different: The pure correlation practice groups associated events that pertain to the same occurrence founded on a previous knowledge: An earlier vulnerability scan, generalized assumptions and/or the technician ability to define significant matching parameters²⁶. On the other hand, Data Mining *discovers* patterns or relationships from large amount of traffic without a prior user knowledge²⁷ using a wider scenario that can include the correlation capabilities.

Conclusion And Future Directions

Intrusion Detection and Data Mining techniques integration process are in a phase of interesting technical and research development, where they have been combined to obtain tested performance improvements over current IDS commercial products, with special features as:

- Data Mining based IDS improves the feature extraction and meaningful information from large amount of raw data, getting more accurate models results from applying DM techniques and correlation algorithms. These practices reduce notoriously the overload analysis to human operators and the associated issues.
- Potential for predictive analysis capability over suspicious activity, enabling defensive action before a severe injure occurs or the system can be totally compromised while reduces the latent associated cost.
- The data mining adaptability for specific environments allows recognize individual trends more than generalize over attack models making appropriated the statistical recognition of new or hidden malicious activity variants of known trends.
- Specific applications as the alarms management show significant improvements in the IDS performance and accuracy.

The complexity to set up this kind of systems has faced with new challenges to the investigators toward scalable and incremental deployments. The goal: to ensure accuracy with efficient time and resources investments specifically improving over data selection, data preparation, data quality (from all the information sources), result precision and the associated computing costs.

In addition most of the efforts are focused to improve its approach to real time performance detection and attacks analysis capability in distributed environments improving the answer time and as well as including specific items as IDS load

balancing, multi-level filtering and the use of complementary techniques as Data Fusion²⁸. The goal: Refine its results and reduce its false positive rate to take advantage of its potential as trusted source able to action with smart responses in real environments.

Finally the promising development in this area has revealed significant advance for a new IDS generation that represents Data Mining based IDS: It enables the potential to move from a reactive security scheme toward one more preventive against the criminal activity.

References

¹ Richard, Power. "2002 CSI / FBI Computer Crime and Security Survey Vol. VIII" Computer Security Institute Spring 2002.

URL: http://www.gocsi.com/db_area/pdfs/fbi/FBI2002.pdf (Jun 10th 2003)

² IntruVert Networks. "Top 10 requirements for next generation IDS" Sept 2002

URL: <http://www.securesynergy.com/library/whitepapers/pdf/intruvert-01.pdf> (Jun 10th 2003)

³ R. Coolen. "RTO Technical Report 49 -Intrusion Detection: Generics and State-of-the-Art –Research & technology Organization". Jan 2002

URL: <http://www.tno.nl/instit/fel/ts/resources/rto-tr-049-ids.pdf> (Jun 10th 2003)

⁴ Ranum, Marcus. "False Positives: A User's Guide to Making Sense of IDS Alarms" ICSA Labs. Feb 2003

URL:

<http://www.icsalabs.com/html/communities/ids/whitepaper/FalsePositives.pdf>

(Jun 10th 2003)

⁵ Norton, Peter. Peter Norton's Network Security Fundamental. SAMS, Indiana 2000 Pg 198

⁶ Stiennon, Richard. Matthew, Easley. Intrusion Prevention Will Replace Intrusion Detection. Gartner Research. Aug 2002

⁷ Power, Richard. "2002 CSI / FBI Computer Crime and Security Survey Vol. VIII" Computer Security Institute. Spring 2002.

URL: http://www.gocsi.com/db_area/pdfs/fbi/FBI2002.pdf (Jun 10th 2003)

⁸ Mulhall, Peter. Zaritsky, Raul Michael . "Knowledge Discovery In Databases - NCSA ".Sept 1999.

URL: <http://cilt.berkeley.edu/seminar/2000/Zaritsky/tsld001.htm> (Jun 10th 2003)

⁹ Abraham, Tomas. "IDDM: Intrusion Detection Using Data Mining". Department of Defense- DSTO Electronics and Surveillance Research Laboratory. May 2001.

URL: <http://www.dsto.defence.gov.au/corporate/reports/DSTO-GD-0286.pdf> (Jun

10th 2003)

¹⁰ Schultz, Matthew. Zadok, Eres. Stolfo, Salvatore. "Data Mining Methods for detection of New Malicious Executables". Columbia University.

URL: <http://www.cs.cornell.edu/people/schultz/papers/ieeesp01.pdf>

(Jun 10th 2003)

¹¹ Julisch, Klaus. "Mining Intrusion Detection Alarms for Actionable Knowledge". IBM Research. 2002

URL: <http://www.zurich.ibm.com/~kju/KDD2002.pdf> (Jun 10th 2003)

¹² Lee, Wenke. "A Data Mining Framework for Building Intrusion Detection Models". Computer Science Department Columbia University. 1999

URL:

http://citeseer.nj.nec.com/cache/papers/cs/18513/http://zSzzSzwww.cs.umbc.edu/SzcadipzSzdocszSzNetworkIntrusionzSzieee_sp99_lee.pdf/lee99data.pdf (Jun 10th 2003)

¹³ Noel, Steven. Charles, Youman. Duminda, Wisekera. . "Modern Intrusion Detection, Data Mining and Degrees of Attack Guilt". Center for Secure Information Systems George Mason University. 2001

URL: <http://www.isse.gmu.edu/~snoel/IDS%20chapter.pdf> (Jun 10th 2003)

¹⁴ Lee, Wenke. "Real Time Data Mining-based Intrusion Detection". Computer Science Department, North Carolina State University. June 2001

URL: <http://www.cs.columbia.edu/~sh553/papers/others/dmids-discex01.pdf> (Jun 10th 2003)

¹⁵ Abraham, Tomas. "IDDM: Intrusion Detection Using Data Mining". Department of Defense- DSTO Electronics and Surveillance Research Laboratory. May 2001.

URL: <http://www.dsto.defence.gov.au/corporate/reports/DSTO-GD-0286.pdf> (Jun 10th 2003)

¹⁶ Lee, Wenke, Park, Christopher. "Automated Intrusion Detection Using NFR: Methods and Experiences". Computer Science Department Columbia University. 1999

URL: http://www.usenix.org/publications/library/proceedings/detection99/full_papers/lee/lee.pdf (Jun 10th 2003)

¹⁷ Bloedorn, Eric. "Data Mining for Network Intrusion Detection: How to Get Started"

URL: http://www.mitre.org/work/tech_papers/tech_papers_01/bloedorn_datamining/bloedorn_datamining.pdf (Jun 10th 2003)

Klaus Julisch. "Mining Intrusion Detection Alarms for Actionable Knowledge" IBM

Research- 2002

URL: <http://www.zurich.ibm.com/~kju/KDD2002.pdf> (Jun 10th 2003)

IBM. "IBM Delivers Autonomic Security Management Software". Sept 2002

URL: <http://www.ibm.com/software/news/n/tkin5fsnj8> (Jun 10th 2003)

²⁰ Broderick, John. "Network intrusion-detection solutions". IDG.2001

URL: http://www.idg.net/crd_detection_16738.html (Jun 10th 2003)

²¹ Intrusion.INC. "Enterprise IDS Management Intrusion SecureNet ". 2001

URL: https://www.intrusion.com/products/downloads/NidsPO_1102.pdf (Jun 10th 2003)

²² Garrison, Mike. Black, Steve. "Managing intrusion detection sensors with Tivoli SecureWay Risk Manager". April 2001

URL: <http://www-106.ibm.com/developerworks/library/it-0401art2/index.html> (Jun 10th 2003)

²³ Hervé, Devar. Andreas, Wespi. "Aggregation and Correlation of Intrusion-Detection alerts". 2000.

URL: http://www.cc.gatech.edu/~wenke/ids-readings/Herve_Debar_IDS_Correlation_Raid01.pdf (Jun 10th 2003)

²⁴ Shipley, Greg. "Dragon claws on the top". Aug 2001

URL: <http://www.networkcomputing.com/1217/1217f2.html> (Jun 10th 2003)

²⁵ ISS. "Real Secure Site Protector Frequently Asked Questions". Version 1.0. 2002.

URL: <http://documents.iss.net/literature/SiteProtector/RSSPFAQ.pdf> (Jun 10th 2003)

²⁶ ISS. "Enhanced Dynamic Threat Protection Via Automated Correlation and Analysis". 2002.

URL: http://documents.iss.net/whitepapers/Dynamic_Correlation.pdf (Jun 10th 2003)

²⁷ VanBelleghem, Dan. "Solving Network Mysteries, Part 2: Auditing and Monitoring Techniques" Dec 2002.

URL: http://searchnetworking.techtarget.com/originalContent/0,289142,sid7_gci803230,00.html (Jun 10th 2003)

²⁸ Silk Road Group. "IDS Data Fusion"

URL: <http://www.silkroad.com/papers/html/ids/node3.html> (Jun 10th 2003)

²⁹ Phung, Mahn. "Data Mining in Intrusion Detection" Oct 2000

URL: http://www.sans.org/resources/idfaq/data_mining.php

(Jun 10th 2003)

³⁰Goeldenitz, Thomas “IDS – Today and Tomorrow” Jan 2002

URL: <http://www.sans.org/rr/paper.php?id=351>

(Jun 10th 2003)

© SANS Institute 2000 - 2005, Author retains full rights.