APPLYING DATA-MINING TECHNIQUES IN HONEYPOT ANALYSIS

Author and co-authors

Namosha Veerasamy	**Pontjho Mokhomoana**	**Johannes Vorster**
CSIR	CSIR	CSIR
nveerasamy@csir.co.za	pmokhomoana@csir.co.za	jvorster@csir.co.za
+27 82 403 9142	+27 82 553 6336	+27 84 219 5713
+27 012 841 3931	+27 012 841 3931	+27 012 841 2258
P O Box 395, Pretoria	P O Box 395, Pretoria	P O Box 395, Pretoria

ABSTRACT

Very little is known about the exact actions executed by a hacker entering a system. Much insight can be gained from following and understanding a hacker's behaviour. It is believed that the more interesting the target the faster the attack will occur. Honeypots are a means of creating an inviting target to lure attackers with the purpose of studying the attackers and their attack patterns. Understanding these attack strategies, patterns and trends can be helpful in determining the vulnerabilities of a system. This requires the system to capture and log large amounts of data which are very difficult to process manually. This process can be time consuming and usually the results are mainly statistical in nature.

This paper proposes the use of a data mining techniques to analyse the data recorded by the honeypot. This data can also be used to train Intrusion Detection Systems (IDS) in identifying attacks. Since the training is based on real data it will better identify and classify attacks than the rule based intrusion IDS's

KEY WORDS

Honeypot, data-mining, intrusion-detection system

APPLYING DATA-MINING TECHNIQUES IN HONEYPOT ANALYSIS

1 INTRODUCTION

Intruders are keen to take advantage of systems to exploit vulnerabilities. Attackers relish in the opportunity to wreak havoc on systems by stealing or manipulating resources to their advantage. Understanding and recognizing when an attack takes place would play a vital role in improving the security of the system and thus offer better protection.

However very little information is available concerning attacks. Often a user simply tries to repair the damage and continue operation of the system. In most cases the data is captured in normal security logs that are too large to study and understand. With busy networks inundated with large amounts of traffic, it is also possible that some of the traffic can be dropped thus resulting in a loss of attack information.

There are many questions surrounding attacks- what intrusive actions were carried out, what is the reason/purpose of the attack, how was it done? Normal computer security does not hold the answers to these questions. It has always been a principle of defence organisations to collect intelligence on the enemy. In this way the enemy can be studied to understand and defend against attacks. To provide good security attack strategies and methods should be understood. Computer security could also benefit by adopting this approach and gather intelligence on attackers. However with standard computer security not containing this information in an easily readable format, understanding attacks can be quite difficult.

2 HONEYPOTS

Honeypots are one such means that can be used to gather information on attackers. Honeypots are not a new technology and have been implemented in the past. Its usefulness lies in its ability to collect information on the attackers. It is an unconventional security system in that it aims to be exploited by attackers. The reason: there is no better way to learn about attacks than to follow it. A honeypot records the activity of the attacker as they exploit/move through the system. The data captured can then be studied to understand the attacker's footsteps. In this way much insight can be gained on attackers' actions, motives and strategies.

A honeypot is a system designed to be appealing in order to attract the attention of attackers. Its goal is to be probed and exploited in some way. The term is derived from the association of a honeypot being an irresistible target due its sweet and pleasant contents, so too will a honeypot system be enticing to attackers. The aim of attracting attackers is to record their activity so that it can be studied later.

Honeypots are different from other security logs in that they only capture what comes to them. In this way honeypots are not inundated with volumes on network traffic.

It is often a problem to separate attack data from large amounts of normal traffic from a network. This makes honeypots also relatively smaller in size than normal security logs. For example, the Honeynet project [5] generated 1-5 MB of data. A previous honeypot implementation generated between 5-9 MB data a day - much less when compared to the gigs of data produced by other network monitoring tools[5]. However understanding these logs still requires significant effort. One such technique that can be utilised is data mining.

Honeypots can broadly be divided into two categories: production honeypots and research honeypots. A production honeypot serves to mainly improve the security of an organisation by helping to detect intrusions. Many large organisations would deploy a production honeypot to help

protect their networks. It forms a decoy and alerts of the presence of suspicious activity. The second category of honeypots implies the functionality of conducting research. Exploring the activities of the blackhat community will provide insight into their actions. This data can then be studied to identify vulnerabilities and improve security. Research honeypots can be thought of a counter-intelligence medium – gathering information on the bad guys. The information can then be used to protect against threats. [6]

2.1 Good Security

Bruce Schneider describes good security as consisting of prevention, detection, and reaction.[5] Production honeypots providing good security can be considered as follows:

2.1.1 Prevention

Honeypots do little to prevent attacks. Their aim is to attract attackers so that their behaviour can be studied. Preventing attacks consists of using good practices like disabling unnecessary services and installing anti-virus software and regular updates. Best practices are one of the first lines of defence against attacks.

A poorly implemented honeypot could become a harmful tool in the hands of an experienced attacker. Some baseline security has to be implemented on a honeypot system to prevent a complete destruction of data that is being collected for the research purposes.

A honeypot can also be used to serve as a decoy in a larger system. In this way a honeypot detracts the attention of the attacker and prevents attacks to more crucial machines in a system. The attackers spend time exploring the honeypot instead of the actual useful machines. The honeypot can also alert of the presence of attackers and prevent attacks to these more important machines.

2.1.2 Detection

Honeypots can make a significant contribution in the area of attack detection. Intrusion detection systems can pick up too many false alarms or fail to detect actual attacks. Honeypots perform detection as a result of system activity and not signatures. However, studying data collected from honeypots can also assist in determining attack signatures that can be implemented in an intrusion detection system.

2.1.3 Reaction

A honeypot is useful in large system environments as the data collected can be studied later. For example a critical device may be attacked but it is not feasible to remove it from the system and perform an investigation. The compromised machine is repaired and is required to continue operation. A honeypot version of the crucial device would also be attacked and can then easily be removed and examined without affecting the functioning of the main system. The outcome of the analyses should provide some insight into what actions were executed on the critical devices and whether further repair processes are required.

Often after a breach of security the operation of the system cannot be halted. The compromised system can then become so clouded with other traffic it is difficult to retrieve and analyse the attack data. A honeypot provides a suitable solution by having the ability to be removed without influencing the operation of the system.

The aim of this honeypot setup will therefore be to play a detection role at first by collecting the attack data. Thereafter the data collected will be studied to determine attack patterns. This will serve as input for the signature training of an Intrusion Detection System. The system will now be able to react and alert to the presence of certain intrusions thus providing a reactive role. Once deployed as part of a larger network the Honeypot/Intrusion Detection System collaboration will prevent the damage caused by attackers by distracting them with the honeypot and signalling their presence earlier before damage can be done to more vital aspects of the system.

3 PREVIOUS IMPLEMETNATION

Valuable lessons were gained from a previously implemented honeypot. Several test setups were required to achieve a relatively stable operation. With the use of commercial software a basic understanding of the functioning and capabilities is required. Honeypots are meant to be probed and exploited and thus the system is prone to becoming unstable from time to time. The honeypot machines have to be hardened to a certain degree and some security has to be implemented to ensure the system does not fall over as soon as it is operational. (For example installing anti-virus software.) The probability of the machine immediately being detected/scanned after being placed on the open Internet is quite high. Attack activity on the previously implemented honeypot was often found to occur minutes are the system was started up. In addition the machines need to be appropriately networked in order for the logging programs to record the system activity. Placement of hubs and firewalls have to be taken into consideration to allow transparency of the traffic for logging and detection programs, as well offer some form of protection to allow the system to continue operation and not be completely taken down by an attacker.

The research honeypot that was set up as part of a research task at the Council for Scientific and Industrial Research (CSIR- research council in South Africa) collected a substantial amount of attacks in the previous year. These were then studied for statistical and trend patterns. Much insight was gained into the frequency and sources of attacks. The details of the setup were documented in a report entitled "*Honeypot Research Report 2005*" and a discussion of some of the results follows. Several findings concerning time-line trends were detected. The majority of attacks took place in the afternoon period. Attack activity would pick up around tea-time, die down for a while and then resume around lunch- time and peak during the afternoon periods. Two conclusions could be drawn from this- the activity could be accounted for by school children after having returned home probed the system or normal users using their time and resources to search for vulnerable machines at a time frame that usually correlates to normal office hours.[3] A general pattern found was that the majority of attackers were South African based.[3] This limited the scope of the search as Telkom is the service provider and further information could not be easily obtained. However some international attacks were investigated and interesting sources were found. For example a substantial amount of attacks originated from two Telecom Companies in China and Korea. The attacks originating from these two sources tried to exploit SMTP vulnerabilities.[3] Asia is notorious for its spamming practices and this form of attack activity probably arose due to the search of vulnerable mail servers to exploit and thus use for the relaying of spam messages. Information regarding how a mail server is exploited can be investigated in the new task. Ways of detecting and thus protecting against such intrusions can be determined. After discovering a few trends and patterns it is hoped that the additional research proposed will explore other areas of attack study. It is therefore the aim of this task to delve further into the world of the blackhat community and understand how and why attacks are carried out.

Vast amounts of logs can be produced from packet capturing programs like Ethereal. Commercial honeypot software also produces logs for the attacks. It would not be feasible to manually group related traffic based on IP addresses or other criteria. It would also be very tedious to study every log to determine strategy and pattern. The amount of logs generated from the previous honeypot setup demonstrated the need for some form of automatic grouping and studying of the logs generated. It is therefore the aim that data mining techniques can be utilised to facilitate this process.

An initial requirement is that the data will need some pre-processing. The pre-processing of the honeypot data will involve grouping related data together. With the use of data-mining techniques the pre-processing will be facilitated. A decision system will determine whether a sequence of activity should be classified as an attack. Specific sequences of traffic (all packets relating to a sequence of activity) will be fed through the decision system (for example a tree) to check whether the traffic can be categorized as an attack. Through the use of tree logic attack

confirmation/denial will take place. If the behaviour is deemed an attack the pattern detected can be incorporated into the IDS to detect future occurrences of the attack.

Honeypots like IDSs provide a security role. Both are concerned with attacks on systems and the two are quite closely related. A collaborative effort can be conducted by using the research from studying the attack patterns of honeypot data as input for training the IDS and identifying attack signatures.

4 DATA MINING

Data mining is the process of analysing data with the purpose of identifying patterns or relationships. In other words you mine data order to extract previously unknown and potentially useful information from the data. Applications of data mining range from the analysis of medical data to profiling customers. These outputs are then used as a basis for making future decisions. The techniques used to perform the analysis include, but not limited to Neural Networks, Rough Sets, Fuzzy Logic and Decision Trees.

We propose using decision trees to perform the analysis of the data. A decision tree is a tree structure where the nodes represent a test and each branch represents a possible attribute. The terminal nodes represent the decision. Decision trees were chosen in this research because they are simple to understand and interpret, are robust to noisy data and perform very well given large data sets.

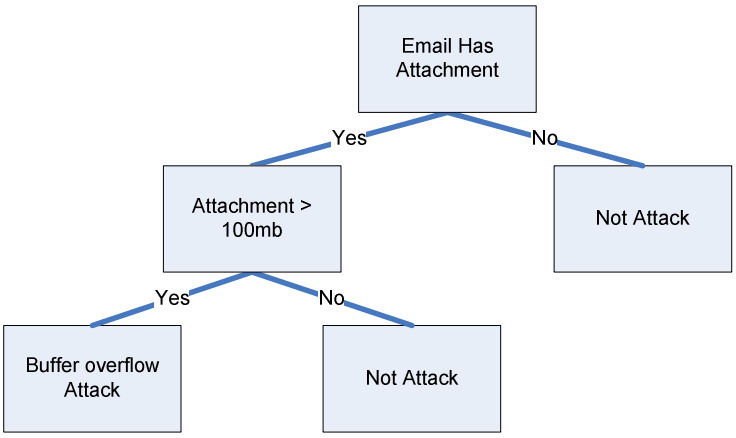

Figure 1: Sample Decision Tree

The first step in the data mining process is the collection of data. The data used in this research will be obtained from a honeypot that will be set up. Winpcap will be used to capture the traffic which will then be archived for later analysis. The next step is preparing the data. The data will be clustered using criteria like session identifier or source/destination address. These clusters will then be labelled and have meta-data such as size, address or protocol attached to them. These clusters are then divided into two sets, the training set and the validation set. The clusters from the training set are fed though the rule induction algorithm to generate the rules. The accuracy of the rules is determined by how well it classifies data from the training set.

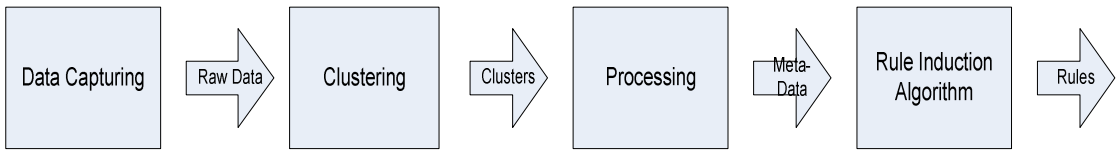

Figure 2: Data Mining Process

We hope to develop signatures for different attacks. The attacks to the honeypot will be simulated so that we can associate an attack type with each cluster. These signatures can then be used by IDSs to identify attacks.

5 IDS

The aim of an Intrusion Detection System is to identify unwelcome, ill-intentioned traffic. This would include of exploitation of vulnerabilities, unauthorised access to information, unwarranted data modification, monopolisation of a system, privilege escalation or installation of worms, viruses or Trojans.

An IDS often contains a form of:

- Sensing
- A display of alerts and detection control
- A database to record the activity and produce a security alert if a match of an attack signature matches the rule set.

IDS can be implemented in several ways and an overview was provided in the Wikipedia encyclopaedia article "*Intrusion Detectio*n System"[7]:

5.1 Signature-based and Anomaly Detection

A signature-based system operates by identifying malicious behaviour through the process of looking for specific patterns in network traffic that match a signature set of attack patterns. In this way "known" attacks are detected. However it is also possible that a new attack can be detected if the new attack shares some features of a known attack.

The signature-based IDS compares the traffic and activity to a database of attack signatures. The IDS will search for attacks that have already been recorded in the database. A signature-based IDS is similar to virus-protection software in that its worth lies in its database of signatures. Regular updates are essential to keep up with new exploits.

An anomaly based IDS monitors the traffic and generates security alerts when "abnormal" behaviour is detected. This is usual performed through the process of self-learning.

Anomaly detection is implemented by having the system administrator set values relating to normal load activity, protocols, average packet size, etc. The anomaly based IDS therefore compares the network segments to the normal indicator levels and watches for irregularities.

5.2 Network- based and Host-based

In a network based IDS the sense mechanisms would be placed at throttle positions in the network to observe the overall traffic and monitors many hosts. Placement is usually at the network periphery, or linked to a hub or a switch. Snort is such an example

Host-based IDS usually have an agent on a host that detects intrusions for that specific host. It functions by analysing systems calls, logs, and modification (like password files) and various other conditions of the host.

Hybrid systems are a combination of the two. The individual host data is added to the overall network information to develop a wide-ranging analysis of the network.

In this task, it is aimed to continue work on a signature based IDS. Using the data from the honeypot it is hoped that substantial attack patterns can be detected to then train the IDS to identify and alert of the presence of these kind of intrusions. In this way more attacks can be sensed, alerted of and the database of attacks can be updated as more attack patterns are found to occur.

6 APPROACH

A honeypot system will be set up which will be the medium of data collection. A honeypot setup requires substantial design and understanding of network technologies as well as the necessary software and different configuration options. Machine, IDS and the logging capability placement has to be considered to achieve data capture as well incorporation of sufficient security measures to offer some level of concealment of the data capturing and attack detection mechanisms.

After considering the various factors that will affect the setup, testing is still required to achieve a successful environment. Much experimentation has to be conducted to ensure the data is accurately and timeously being collected as well as disguise the honeypot and IDS operation. Much testing will be required to achieve these requirements of the system.

In the previous honeypot setup experimentation was required with regards to the placement of the hub for data capturing transparency. This will also be the case as further components in the form of the IDS and additional machines to form a network will also have to be incorporated. A preliminary conceptual representation of the setup is shown in Figure 3.

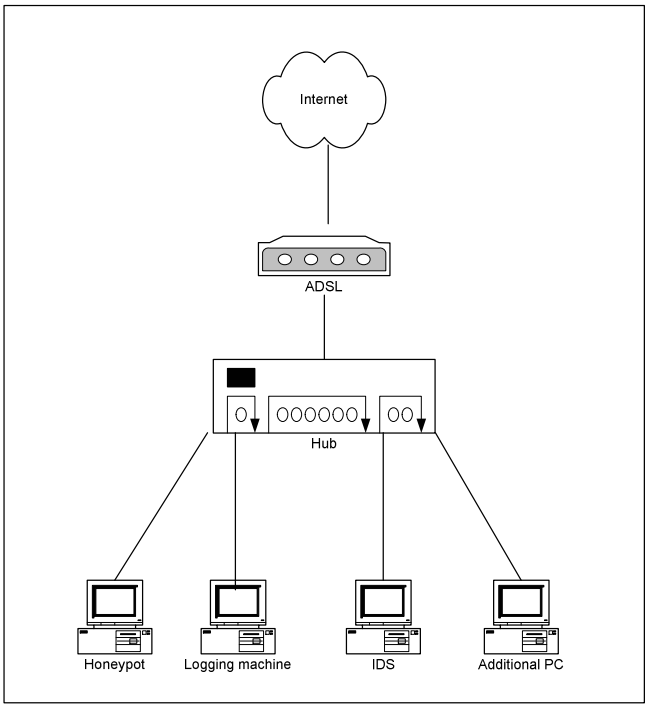

Figure 3: Preliminary Conceptual Honeypot Setup

A honeypot machine, IDS, logging machine and additional machines to form a network are required. The networked setup is connected to an open Internet connection (for example an ADSL line).Further considerations to the setup will have to be made to ensure the IDS is disguised, the attack data is recorded and the system is functional. This will involve a trial-and-error approach as was used in the previous honeypot implementation. The honeypot machine can run commercial honeypot software or be more appealing to attackers, can host a web, file or mail server. The logging machine will run a packet-capturing program like Ethereal or Commview to capture the traffic. Commercial honeypot software also collect logs, though these might be limited in terms of useful information and might not contain the actual traffic but merely a description of the attack.

Commercial honeypots software probably have their own signature based detection of attacks. The logs collected from commercial software programs also might only contain the details of the attack and not the actual traffic. Another consideration is therefore to investigate the link between attacks detected by the honeypot software and the logging program data.

Logging programs record the traffic on the system. The attack data will be contained in this traffic capture. This data will require some pre-processing before being mined to determine attack occurrences and patterns. This pre-processing will consist of linking related data and grouping packets that should flow in sequence of each other and looking for preliminary patterns (an aim of data-mining).

One approach of studying the data is to simulate the attacks and try and determine a pattern in the data. This is helpful as certain characteristics of the attack will be known and this will be searched for in the traffic. This assists with the task of studying and understanding the traffic and classifying certain patterns/behaviour as an attack. In the previous implementation of the honeypot, simulated attacks were useful in testing the functionality and confirmation of attack data capture. Simulated attacks will therefore be a worthwhile consideration.

The overall approach will be to use scripts and data-mining techniques to sort and group related traffic, look for basic patterns and thereafter feed the sorted data through the decision system to detect attacks based on devised criteria. If significant attack patterns are found to occur, this will be incorporated into the design of an IDS to be able to pick up the attacks previously recorded/simulated.

7 CONCLUSION

Understanding how system attacks are carried out will be useful in identifying them and detecting them in the future. Honeypots are a useful technology in the capture of attack data. Due to the magnitude of logs a technique like data mining is proposed to pre-process and sort data. Thereafter a decision system will determine whether the traffic can be classified as an attack. The analysis carried out on the captured data will serve as helpful input in the training of IDS systems in the future. This paper describes lessons learnt from a previous honeypot implementation and explains a research project in progress.

8 REFERENCES

[1] G Kurtz, S McClure & J Scambray, *Hacking Exposed: Network Security Secrets and Solutions*, Osborne/McGraw Hill, 2001.

[2] M Dixon, *An Overview of Document Mining Technologies*, available online from http://www.geocities.com/ResearchTriangle/Thinktank/1997/mark/writings/dm.html, last accessed 2006/04/22

[3] L J Kriel, N Veerasamy & J S Vorster, *Honeypot Research Report 2005*, CSIR, 2006.

[4] S Northcutt, J Novak, *Network Intrusion Detection*, ISBN 0-73571-265-4, New Riders Publishers, Sept 2002.

[5] L. Spitzner, *Honeypots: Definitions and value of honeypots*, tech. rep., Honeynet, 2003. http://www.tracking-hackers.com/papers/honeypots.html.

[6] L Spitzner, *Honeypots: Tracking Hackers*, ISBN 0-321-10895-7, Addison- Weasley, December 2002.

[7] *Intrusion Detection System*, available online from http:// en.wikipedia.org/wiki/Intrusion_detection_system, last accessed 2006/03/10.

[8] N Provos, *Virtual Honeypot Framework*, available online from http://niels.xtdnet.nl/papers/honeyd.pdf, last accessed 2006/04/22.