Analysis of Attacks Using a Honeypot

Gary Kelly and Diane Gan*

CSAFE
University of Greenwich
* D.Gan@gre.ac.uk

Abstract. A Honeypot is a software based security device, deployed to attract hackers by displaying services and open ports which are potentially vulnerable. While the attackers are diverted, their activities can then be monitored and analysed to identify current attack methods and trends. A low-interaction Honeypot called Dionaea was chosen for this project because it can simulate services while preventing an attacker from gaining full control. Results were collected over the six week period of the experiment. The logged information of the observed attacks was analysed and compared with current vulnerabilities, the locations where the attacks were originating from and the time of day at the originating site. A profile of individual attackers can then be built to gain an insight into the current attack trends in order to improve network defences.

1 Introduction

Honeypots are being used increasingly by organisations to detect the presence of attackers [1]. This means that the defenders can keep the attacker ring fenced where they can do little harm, and learn more about the tactics that are currently being deployed in order to fine tune their defences appropriately. There are many advantages of Honeypots because of their simple concept that gives them powerful strengths. However, a Honeypot does not replace existing security technologies but can work alongside them, tracking and capturing activity occurring on the system where it is deployed. However, it will only capture activity that is directed at the Honeypot itself. The Honeypot can also be at risk of sabotage and could be used to attack other connected systems [2].

Honeypots can be classified in one of two ways depending on their deployment as Production Honeypots or Research Honeypots [3]. Production Honeypots are easy to use, capture only a limited amount of information and are primarily used by companies and corporations. Placed outside of the production network, Production Honeypots are used in conjunction with other production servers in order to improve the current existing level of security but give less information about the attackers and the attacks that they mount [3]. Production Honeypots can also be further classified as Low-Interaction Honeypots and High-Interaction Honeypots. Research Honeypots are deployed to offer information into the motives and techniques of the Black Hat community. These are used to research the current threats and to provide information to the organisation about the different avenues of protection against these threats. [3], [4]

Low-interaction Honeypots simulate services which cannot be exploited by an attacker, as they are limited in functionality. However, they are very useful for gathering information at a higher level, such as when analysing worm activity or network probes [4][5]. Examples of Low-interaction Honeypots include Dionaea, Specter, Honeyd and KFSensor [5][2].

Low-Interaction Honeypot	**High-Interaction Honeypot**
Emulates OSs and services	No Emulation; real OS & services provided
Simple and easy to install and deploy.	Can be complex to install and deploy
Minimal risk as the emulated services control what attackers can and cannot do.	Increased risk as attackers are provided with real operating systems to interact with.
Captures limited information which is mainly transactional data and very limited interaction.	Can capture far more information including new tools, communication and attacker keystrokes

Table 1. Comparison of Low And High Interaction Honeypots [2]

High-interaction Honeypots run real services and servers. There is obviously a danger that an attacker could use a high-interaction Honeypot to attack these services, which requires additional technologies to be implemented to prevent this [2]. Examples of High-Interaction Honeypots include Symantec Decoy Server and Honeynets [5][2]. Table 1 compares the main differences between these two types of Honeypot.

The aim of this work was to deploy a Honeypot in collaboration with a penetration testing company located in nearby Canary Wharf. The Honeypot deployed for this work was Dionaea which was set up to present a number of vulnerable services and open ports to the Internet with the intention of attracting attacks for further analysis. The information logged by the Honeypot would be used to identify any correlation between current vulnerabilities and the analysed information. The rest of this paper is organised as follows. Section 2 discusses the deployment of the Honeypot. Section 3 presents the analysis of the Honeypot logs. The Conclusion is given in section 4.

2 Deployment of the Honeypot

Dionaea was chosen as the Honeypot to be implemented for this work. Created for the Summer of Code 2009 by Markus Kötter, Dionaea was the successor to a previously used Honeypot named Nepenthes [6]. Dionaea is a low-interaction Honeypot, which was written in C with an interface based on Python. It supports both IPv6 and Transport Layer Security (TLS) and has real time notifications using XMPP (Extensible Messaging and Presence Protocol). An SQLite 3 database logs the information on each attack and can also produce graphical statistics.

Ubuntu version 11.04 (Natty Narwhal) was the platform of choice for deploying the Honeypot and the following software was installed :

- OpenSSH Server
- DNS Server
- LAMP (Linux, Apache, MySQL, PHP) Server
- Samba File Server

These programs were chosen because the OpenSSH server allows remote connections to the server, which in turn will allow the remote administration of the server. The DNS Server was installed because attackers need to think they are intruding on a LAN. A LAMP server was also installed with a MySQL Database running on port 3306. When all of these programs are used together, they support Web application servers. The combination of both the OpenSSH software and the open port 22 made this a tempting target. A Samba File Server was installed because it can interoperate with a Windows Server Domain, which makes it appear as if more than one server was present on the network. The Dionaea configuration file was amended to install a passive OS fingerprinting tool called p0f, to improve the format of the log files. The Honeypot was run in twelve hour blocks over a period of six weeks and the software p0f was used to convert the logged information into a readable format for analysis.

3 Results

Each individual IP address was logged every time it accessed the Honeypot. Of course some only appeared once, but there were a number of IP addresses which launched multiple attacks, either on the same day or over a number of days. The information from the log files was copied to a spreadsheet for further analysis.

3.1 Overview of the Attacks

The log files revealed that a large number of attackers used the Linux operating system to launch their attacks and connection requests. Only 3% of attacks originated from Windows machines. This left 30% that were unknown because the Honeypot did not recognize their signatures. There were an even smaller number of attempts to remotely connect to the server in order to remotely administrate it from the Sun Solaris System.

Compensating for the time zone of the attacker, it was found that over a typical six day period the majority of attacks occurred between 8 am and 6 pm. Midnight to 8am had the least number of attacks. It can be inferred that either these attackers were unemployed or they do this for a living, i.e. are professional hackers. It was also noted that the Honeypot was most active on a Saturday.

To identify the country of origin a geolocator called iplocation.net was used, which utilizes three different IP geolocators, for greater accuracy. The results were also cross-checked using Traceroute. There were found to be over 100 unique IP addresses logged and a number attacked the Honeypot more than once. The five countries with the highest attack volumes are shown in Fig. 1.

There were a number of daily attacks originating in China with at least one attack per day. These amounted to 33% of all attacks. One IP address that originated in Shanghai, China, was found to have attacked port

22 on three different days. Those originating from the US were only 3% of the total.

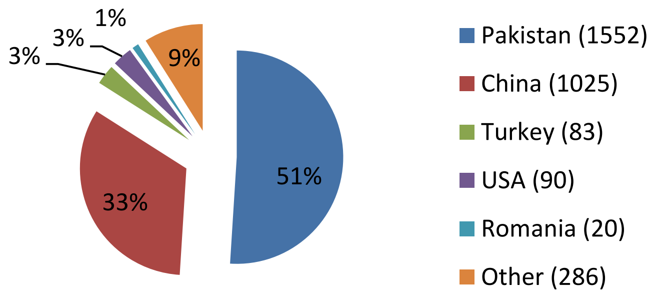

Fig. 1. Five Countries Where the Most Attacks Originated

3.2 Port Vulnerabilities

The three most attacked ports are shown in Table 4. An investigation into these ports and the current vulnerabilities was undertaken to determine if there was a link [7], as discussed in section 3.3.

Service	Port	Number of attacks	Percentage
Secure Shell	22	2774	90.8%
Telnet	23	57	1.9%
Microsoft Active Directory	445	48	1.6%

Table 2. Most Commonly Attacked Ports

Port 22 represented 90.8% of all attacks. A search of current vulnerabilities for OpenSSH server discovered an exploit that allows remotely authenticated users to obtain sensitive information using the auth_parse_options functions that provide debug messages that contain command options. By reading these debug messages, it was possible for this vulnerability to cover the root account because a user account may have no access to any shells or file systems and therefore will have no way to read a particular file in its own home directory [7]. If an attacker exploiting this vulnerability read the debug messages continuously using an automated attack, they could potentially gain unauthorised access to the system. It is possible that the attackers knew of this vulnerability and were attempting to execute a DoS attack on the server in order to gain access while the server was unresponsive, but before it become disabled.

Telnet on Port 23, had the second highest number of attacks. Most current vulnerabilities here were buffer overflows as the attacker attempted to elevate their privileges [7]. Buffer overflow occurs when a program is manipulated into writing to memory blocks outside of the allocated buffer space. Using inputs that exploit code and alter program operation an attacker can gain access to a system. In this case they were trying to crash the server to perform a DoS attack.

The attack on the Microsoft Active Directory (port 445) was surprising because the server OS was a Linux distribution. The latest vulnerabilities here included a DoS attack using a specially crafted query to the Microsoft LDAP (Lightweight Directory Access Protocol) [7]. It is possible that the attacker knew of this vulnerability and was attempting to use it for a DoS attack on the server.

3.3 Analysis of Attacked Ports

Table 2 shows the top ten ports attacked. It lists the service that operates on each port, the number of times that port was attacked and the percentage that each attack contributed to the total number of attacks on the Honeypot.

The results in Table 2 show that the attackers were attempting to identify services that were proprietary to Microsoft. Port 445 (Active Directory) and port 3389 (Microsoft Terminal Server) were both targeted. MS Active Directory is used by Windows Domain Networks, for authentication purposes and for determining directories, policies and services. Microsoft Terminal Server is used to host multiple client sessions on Windows Server Operating Systems. This implied that the attackers were unaware that the server was running a Linux OS. The MySQL Database (port 3306) was also targeted. MySQL Database is part of the LAMP server so this attack was not unexpected.

Telnet (port 23) had the second highest number of attacks. This was attributed to the buffer overflow vulnerability that was present at the time of the experiment, with attackers trying to gain root access.

The largest number of attacks were aimed at port 22 the Secure Shell. On one day there was a very high volume attack (1595 separate loggin attempts) on this port from a single IP address, which was traced to Mirpur, Pakistan using the geolocator. These occured over a very short period of time, indicating an automated attack. Further analysis of the log data showed that the source IP had been up for 7,263 hours (over 300 days). This indicated that a server probably running a Botnet was performing the automated attack. On another day 440 attacks out of a total of 457 attacks occurred against port 22, which accounted for 96.2% of the total for that day. Further investigation found that this attacking IP source had been running for 6,728 hours (around 280 days). Again it was assumed that this could have been a Botnet or an automated attack launched from a server.

Port #	Service	Individual Attacks	Percentage of Total Attacks
22	Secure Shell(SSH)	2774	90.8%
23	Telnet	57	1.9%
445	Microsoft Active Directory	48	1.6%
3389	Microsoft Terminal Server	41	1.3%
4899	Radmin Remote Admin	28	0.9%
135	Microsoft EPMAP	26	0.9%
1433	MSSQL	15	0.5%
3306	MySQL DB System	15	0.5%
5900	VNC Remote Desktop Protocol	15	0.5%
443	Secure HTTP	7	0.2%

Table 3. The Top 10 Ports Attacked

3.4 IP Addresses of Interest

A number of IP addresses were investigated further, because they launched a significant number of attacks or because they were found to have some anomaly regarding the source location. Table 3, lists these IP addresses, the source location and the reason for further analysis.

The IP address, 27.54.120.3 was traced in Mirpur, Pakistan. As these attacks occured in the same day and in a very short time frame (3 second intervals) it can be assumed that this was an automated attack.

The IP address 103.6.220.215 was tracked to Shanghai, China using the geolocator iplocation.net. Attacks were launched from this IP address on three days. The first was a single probe, followed by two larger attacks, see Table 3.

IP Address	Location	Reason For Further Analysis
27.54.120.3	Mirpur, Pakistan	1595 attacks on one day
103.6.220.215	Shanghai, China	1, 440 and 439 attacks over 3 days

Table 4. IP Addresses Needing Further Investigation

4 Conclusion

The deployment of a Honeypot can be beneficial in determining current attack strategies and probes being used by hackers due to the information collected. Using the Dionaea Honeypot, attacks were identified on a number of ports, which were compared to current vulnerabilities, although these were not the only ports attacked. Many attackers were scanning for characteristic Microsoft ports. The most popular port attacked was port 22, as a compromise here can have the biggest reward for the attacker. Detailed analysis of the logged information found that the attackers were based all over the world, including Eastern Europe, Russia, South America, Middle East, Greece, China, India and Pakistan. The time of day at the source was interesting as the majority of attacks occurred during the working day.

5 References

1. Robert Lemos, 5 Reasons Every Company Should Have A Honeypot, 1st October 2013, Accessed 23 March 2014,http://www.darkreading.com/advanced-threats/5-reasons-every-company-should-have-a-ho/240162106
2. Spitzner, Lance. “Honeypots: Definitions and Value of Honeypots”, May 2003, accessed: November 2012, URL: http://www.tracking-hackers.com/papers/honeypots.html
3. Anonymous, “Honeypot (Computing)”, Date Accessed: October 2012, http://en.wikipedia.org/wiki/Honeypot_(computing)
4. Almutairi, Abdulrazzaq “Survey of High Interaction Honeypot Tools: Merits and Shortcomings”, June 2012, Date Accessed: October 2012 http://www.cms.livjm.ac.uk/pgnet2012/Proceedings/Papers/1569604821.pdf
5. Honeyd, “Honeypot Background”, Date Accessed: October 2012, URL: http://www.honeyd.org/background.php
6. anon. (2009). The Honeynet Project, Google Summer of Code 2009. Available: http://www.honeynet.org/gsoc2009. Last accessed 6th 4 2014.
7. Common Vulnerabilities And Exposures. “CVE-2012-0814” , 16th February 2012 , accessed: 29th April 2013 http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2012-0814