

An Overview of ADSL Homed Nepenthes Honeypots In Western Australia

Craig Valli and Aaron Wooten
School of Computer and Information Science
Edith Cowan University
c.valli@ecu.edu.au
awooten@student.ecu.edu.au

Abstract

This paper outlines initial analysis from research in progress into ADSL homed Nepenthes honeypots. One of the Nepenthes honeypots prime objective in this research was the collection of malware for analysis and dissection. A further objective is the analysis of risks that are circulating within ISP networks in Western Australian. What differentiates Nepenthes from many traditional honeypot designs it that is has been engineered from a distributed network philosophy. The program allows distribution of results across a network of sensors and subsequent aggregation of malware statistics readily within a large network environment.

Keywords: *honeypot, Nepenthes, malware*

INTRODUCTION

Nepenthes (Wicherski, 2007) is a malware collection program that emulates vulnerability in Microsoft Windows operating systems. The Nepenthes system typically runs on a Linux or UNIX based system providing honeypot capabilities. Atypically honeypots are focused on tracking interactions between the attacker and the victim machine, Nepenthes however, is focused on the collection of malware and consequently has a significant change in *modus operandi*. Nepenthes does have significant logging capability and can be used alongside established honeypots such as honeyd to track interactions, however, one of its primary purposes is the collection of malware for analysis. It works by emulating known vulnerabilities such as MS03-26 (Microsoft, 2003) such that it will receive a malicious payload should it be available from the attacking entity. The emulation of the target vulnerability itself it should be noted is not a complete replication of malware signature. The emulation is sufficiently convincing to deceive the attacking entity into believing they have successfully produced a compromise in the host and are therefore confident enough to transmit the malcode to the Nepenthes server.

What further differentiates Nepenthes from traditional honeypot designs it that is has been engineered from a distributed network philosophy. Initial modules are created in the program to allow distribution of results across a network of sensors. This allows for the aggregation of malware statistics readily within a large network environment. Furthermore, Nepenthes has significant SQL logging capabilities allowing logging to popular SQL servers such as MySQL and PostgreSQL database systems for later aggregation and analysis with other tools one such implementation is Surfnet.

Surfnet (ids.surfnet.nl, 2007) is a series of programs used to aggregate and analyse the data that is trapped within a Nepenthes honeypot architecture. Surfnet is a combination of open source tools to produce a fully distributed sensor network and logging system. The use of VPN technologies (OpenVPN) allows for the use of VPN tunnels to transmit data from the sensors to the central logging computer. The model for which is illustrated in Figure 1




Figure 1: Surfnet Design (*ids.surfnet.nl, 2007*)

The Surfnet model as can be seen clearly from Figure 1 allows for the easy separation of the honeypot from the logging. Other honeypot systems make it difficult or often impossible to place logging and monitoring readily on another system other than the one that is running the physical honeypot. The sensors in this design can be a CD-ROM or USB based implementation of the system this allows for very cheap resilient sensors. The sensors are designed to be updated across the wire via the use of the Subversion revision control system to update configurations and binaries.

The logging server uses a variety of technologies and techniques to readily analyse the data that is drawn in from the honeypot activity. This activity is illustrated in Figure 2




Figure 2: Nepenthes system

The use of Google maps and geographic locational services to produce a worldwide map of attackers is just one of the readily available analysis tools in the Surfnet system. The web interface also has the ability to search by each of the various attributes of the captured data set for example by attacking IP number, network sensor or malware file.

WHERE IS THE VALUE IN THIS FORM OF HONEYPOT?

There are numerous honeypots that are designed to map the interaction between attacker and victim machine. There is sufficient evidence produced that these systems are in of themselves an invaluable tool for tracking malicious activity on a network. One of the attendant problems with these types of systems is the ongoing maintenance and analysis required for them to be effective. These systems typically presume that there is some human factor in the interaction or attack. However, this paper argues that this type of attack is becoming somewhat of an arcane way to attack and ultimately compromise systems and that there is a shift in the motivation for compromise of machines.

The increasing use of commoditised and fully automated attack tools such as nmap, nessus and others is making the attack of systems a relatively trivial task. Furthermore, the increasing customisation and simplification of malware production through the use of tools such as the Metasploit framework (www.metasploit.com, 2007) is presenting significant challenges for the detection and remediation of malicious code. These frameworks will generate exploit that have various and often unique signatures many of which evade detection by heuristic analysers.

Research into the interactions may allow the production of attack signatures which by their very nature will allow the alerting of systems that an attack is occurring but often ultimately may not stop or prevent the compromise. New or unknown (zero day) compromises of systems are often unavoidable and can take considerable time to mitigate. This is due to a variety of issues which could be the nature of the attack being based in a protocol such as the 802.11b exploit of WEP. The fact that the software vendor may not have a fix for the particular compromise let alone be aware of it is an increasingly common phenomenon.

The stereotyped notion of a hacker's profile being a solitary attacker bent on destruction or simple system penetration is dated as much as it is ignorant. The modern hacker or *üebercracker* (Venema & Farmer, 1993) is a reality and is now network enabled and understands programming, exploit and automata of same. There is mounting evidence to suggest that the *modus operandi* of attackers is changing from a merely curiosity or personal fame perspective to one increasingly motivated by personal gain or profit through the concept of covert ownership of third-party computers. The concept of "oWning" a computer whereby a system is compromised for later purposeful use is at odds with the conventional wisdom on hackers motivations which is the defacement or destruction of a system for fame. It moves the malicious intentions from a predatory one to parasitic one based on a profit motive. Profit here does not necessarily mean an exchange of fungible goods but the availability of botnets for use in denial of service on a fee basis on the Internet is just one example of this growing trend of "ownership" with a profit imperative rather than a destructive one. There have been documented cases where these botnets have been used to create denials of service on victims with subsequent demand for the money to desist from the denial of service. In the case of online casinos or other businesses that require and rely on service availability as a revenue stream, payment of these demands is a cheaper option than continued denial of service. In addition now organised crime syndicates and enterprising hackers now produce commoditised and customised malware. These off the shelf malware can be purchased on-line and have full support infrastructures rivalling or exceeding service levels given by legitimate vendors to make your malware dollar go further. Hence lessening the need for expert or specialist IT knowledge to compromise or "oWn" a network of computers or enterprise.

The major and reducible threat is the "ownership" of the victim machine by a malicious entity as a result of executing malicious code on the victim machine that allows control of the computing resource by a third-party. The reduction in threat is achievable by disallowing the execution of the code itself or by disallowing the external communication to the third party which allows for control of the system on which the new malcode has been executed. One way of preventing this is the reduction in threat achieved through the collection and dissection of the offending malcode. The dissection of the malcode can allow mitigation of the threat by disallowing the execution of the code or achievement of the objective which is typically open of a covert channel for control or compromise of valuable data. As a simplistic minimum it can provide a signature for intrusion detection systems to deny download of the malware.

OUTLINE OF NEPENTHES COLLECTION NETWORK

The current network is a network of 5 – 8 computers deployed within the geographical locale of Perth, Western Australia. The sensors are located on various ISP networks all are commodity based ADSL connections from ADSL 512K through to ADSL2 24Mbit connections. These are further classified as being National or Tier 1

these ISPs are major national infrastructure backbone providers, Local Tier 2 – major state based backbone providers and Local Tier 3 who are state based providers who use Tier 2 or Tier 1 for backbone access.

Tier 1	3
Tier 2	4
Tier 3	2

When a piece of malware was successfully downloaded by the Nepenthes engine it was subjected to a range of internal and external tests. Internally the malware was run against a database of updated virus engines this update and subsequent scan of malware is spawned on a daily basis. Externally on arrival all new malware was submitted automatically to the Norman Sandbox and Anubis testing system for decode and detection of exploit. The results for each piece of downloaded malware was then downloaded and stored in the database. Furthermore, all pieces of malware downloaded are submitted to the alliance.mwcollect.org for processing from a global perspective. There is logging of all this activity to standard text log files, SQL and capture of the various outputs to disk.

PRELIMINARY RESULTS

Table 1 outlines the detected connections to the honeypot network. It should be noted that these are connections that the complete network believes to have occurred.

Detected connections	Statistics
Possible malicious attack	70,359
Malicious attack	4,814
Malware offered	4,678
Malware downloaded	949

Table 1 – Detected Connections

The network detected 70359 attacks of which 4814 (6.8%) it believed were malicious attacks. Of these 97.2 % (4678) offered malware to the Nepenthes systems. Of the offered malware only 20.2% (949) resulted in successful download of malware to the system. Of overall possible malicious attacks 70359 these 949 malware downloads only represents 1.35% of all connections made.

Vendor A	601	63.3%
Vendor B	56	5.9%
Vendor C	646	68.1%
Vendor D	70	7.4%

Table 2 - Virus Scanners unidentified malware

Once downloaded the 949 were subjected to analysis by 4 mainstream virus detectors in use. At the last check Vendor B appears to have the best recognition rate of malware with only 56 tagged as Suspicious by the scanners. Vendors A and C do not identify over 60% of the downloaded malware by name or type which is of great concern. At this stage there has been no correlation to these samples with actual reverse engineering and identification of the malware which is underway as a separate research project.

	Address	Total
1	Australia1	1775
2	Australia2	623
3	Japan	412
4	USA 1	387
5	USA 2	364
6	Taiwan	364
7	USA 3	364
8	Spain	364
9	China	364

10	Australia 3	364
----	-------------	-----

Table 3 - Top Attackers by Country

Table 3 indicates the Top 10 attacking IPs by number of connections. It should be noted that Australia1 and Australia2 were on the same ISP networks as 2 of the Nepenthes sensors. Australia3 was from an Eastern States ISP for which all intents and purposes is geographically and topically remote from Australia1 and Australia2.

#	Protocol	Total
1	link	1900
2	creceive	1598
3	ftp	1277
4	tftp	475
5	blink	338

Table 4 - Top 5 download protocols of all sensors

The top 5 protocols as in Table 4 hold no surprises with respect to protocols used to transport malware.

IMPLICATIONS OF RESULTS

The initial results clearly indicate that there are significant and growing problems with the spread of network borne malware. Firstly, some infrastructure providers appear to be blissfully unaware or dilettante in providing a secure as possible network for their customers and the wider Internet community. Seven of the nine ISPs appear to provide a filtered environment for end users while the remaining two ISPs appear content to allow attacks to occur within their networks. These internally focussed attacks are sustained, persistent and prevalent evinced by the top two attacking IPs from Table 3. These ISPs are leaving the way open for possible litigation by allowing these attacks to occur and seemingly providing no mitigation of the malicious behaviour.

The age of some of the exploit code that is still in circulation would point to the fact that many of these older exploits still compromise hosts. MS03-26 is a relatively old exploit for which there is vendor based patching available that stops the exploit. Of further interest is that some of the vulnerable Microsoft platforms for this specific attack are nearing end of life support with Microsoft. Yet these types of attack were high in occurrence in the overall statistics. This trend points to organisations or individuals deploying computers with little or no patching and few protections such as firewalls.

One of the other trends is the speed of released exploit code spread across the networks. There is strong evidence that the release of exploit is often accompanied by a surge in the production of variants or clones of particular malware. What is further exacerbating this problem is the incorporation of code into releases of exploit frameworks such as Metasploit that likewise saw a resultant surge in the number of cloned malware. These surges have significant and profound impacts on signature based systems ability to detect and mitigate threats.

The Surfnet framework allowed for automated scanning of malware against a range of virus analysers and malware detection. Persistent zero day malware was found as a result of virus scanners inability to detect or even quarantine the new malware. The researchers have one piece of malware that remained undetected for a period of 8 months during the research by some popular virus detection engines. There were other pieces of malware that also proved resistant to detection for long windows of time some for as long as 3 months. The statistics in Table 2 see two of the virus detectors leave over 60% of the malware tagged as suspicious being unable or unwilling to make a determination of malware type. This has significant implications for effective mitigation of threats if one was reliant on these virus detection engines for protection. To digress this form of reliance on this form of protection is akin to Rumsfeldian logic "There are known knowns; there are things we know we know. We also know there are known unknowns"(news.bbc.co.uk, 2007)

The framework allowed submission of malware to analysis boxes such as Anubis or Norman Sandbox and macro virus submission sites such as virustotal.com. These malware dissection systems utilise sandboxing and other automation to dissect and analyse suspicious malware submitted to the sites. During the conduct of the current research various malware sites and blogs provided details on how to build malware that avoids detection by the Anubis system.

CONCLUSION

Hackers and corporate criminals are now targeting and using the Internet as a revenue stream. Comprise of network connected computers to create botnets for sale is now almost a mundane occurrence. We need to adapt our modus operandii to mitigating the threat posed by the uebercracker and their legion of automaton attackers utilising their exploit codes or attack tools.

Nepenthes style honeypots used in a widely distributed mode provide a mechanism to gather significant credible attack intelligence and perform environmental scanning beyond traditional research based honeypots. These types of honeypots accept that compromise is not a rarity but an eventuality in the ever changing exploit space. These systems in fact welcome compromise by mimicry of successful exploit and capture of malware for dissection. Honeypot purists may see this as an approach defeatist and one that does not track the elite or discover new exploit. In the limited analysis we have conducted this type of honeypot does successfully trap zero day exploit as was demonstrated in the data collected in this research.

REFERENCES

- ids.surfnet.nl. (2007). SURFids [SURFids]. Retrieved 13 November 2007, from <http://ids.surfnet.nl/wiki/doku.php>
- Microsoft. (2003). Microsoft Security Bulletin MS03-026 - Buffer Overrun In RPC Interface Could Allow Code Execution (823980). Retrieved 6th Feb, 2006, from <http://www.microsoft.com/technet/security/bulletin/MS03-026.mspx>
- news.bbc.co.uk. (2007). BBC NEWS | Americas | Rum remark wins Rumsfeld an award. Retrieved 13 November 2007, from <http://news.bbc.co.uk/2/hi/americas/3254852.stm>
- Venema, W., & Farmer, D. (1993). Improving the Security of your site by breaking into it.
- Wicherski, G. (2007). Nepenthes. Retrieved 13 November 2007, from <http://Nepenthes.mwcollect.org/>
- www.metasploit.com. (2007). The Metasploit Project. Retrieved 13 November 2007, from <http://www.metasploit.com/>

COPYRIGHT

Craig Valli ©2007. The author/s assign Edith Cowan University a non-exclusive license to use this document for personal use provided that the article is used in full and this copyright statement is reproduced. Such documents may be published on the World Wide Web, CD-ROM, in printed form, and on mirror sites on the World Wide Web. The authors also grant a non-exclusive license to ECU to publish this document in full in the Conference Proceedings. Any other usage is prohibited without the express permission of the authors.