

2007

Be Aware with a Honeypot

Stephen Meyer

Ivan Smyth

Mark Cummins

Anthony Keane

Follow this and additional works at: <https://arrow.tudublin.ie/itbj>



Part of the [Digital Communications and Networking Commons](#)

Recommended Citation

Meyer, Stephen; Smyth, Ivan; Cummins, Mark; and Keane, Anthony (2007) "Be Aware with a Honeypot," *The ITB Journal*: Vol. 8: Iss. 2, Article 2.

doi:10.21427/D7WB47

Available at: <https://arrow.tudublin.ie/itbj/vol8/iss2/2>

This Article is brought to you for free and open access by the Ceased publication at ARROW@TU Dublin. It has been accepted for inclusion in The ITB Journal by an authorized administrator of ARROW@TU Dublin. For more information, please contact arrow.admin@tudublin.ie, aisling.coyne@tudublin.ie.



This work is licensed under a [Creative Commons Attribution-NonCommercial-Share Alike 4.0 License](#)

Be aware with a Honeypot
Stephen Meyer, Ivan Smyth,
Mark Cummins and Anthony Keane
Institute of Technology Blanchardstown, Dublin 15, Ireland

Abstract

The Internet has already become a hostile environment for computers, especially when they are directly connected with a public IP address. We have experienced this hostile activity where on an average day; the ITB Honeypot recorded over a thousand reconnaissance attacks seeking unauthorised entry onto our private network. Our Honeypot is a basic PC running Windows XP with no services offered and no activity from users that would generate traffic. The Honeypot is running in a passive state on a stub-network where all inbound and outbound traffic is recorded at the bridging computer to the WAN. We report on the majority of scans and vulnerability attacks that were used and investigate the processes that targeted vulnerable ports and access points on the network.

Keywords: Honeynet, Honeypot, Honeywall, Internet monitoring, Cyber attacks

Introduction

It is generally accepted that the average time for an unprotected computer to be compromised on the Internet now less than two hours. Here we investigate this claim by recording uninvited network activity implemented against our computers that are connected directly to the Internet and constitute our Honeynet [1]. As a Honeynet is an unadvertised network and does not run applications that initiate Internet traffic, then all traffic on a Honeynet is considered malicious and goes through a data control firewall that tracks inbound and outbound connections and an intrusion protection system (IPS) to prevent any compromised Honeypots from being used to initiate attacks by dropping or modifying malicious traffic originating from them.

Vulnerabilities in Networks

In general, computer networks are composed of devices, applications and protocols. The typical devices are switches, routers, servers and client computers. The applications are network and client operation systems, web and email services and many other application services that vary depending on the business using the network. The communications between devices and applications use standard well known protocols [2], many of which have little or no security ability built-in to them. Individually and together, these component parts of the network provide a wide ranging array of weak points (vulnerabilities) that hackers probe and attack in order to gain access and eventually take over computers on the network [3].

Typical points of attack in a network are:

1. Poor configuration of router access controls lists that allow leakage through ICMP, IP, NetBIOS and can lead to unauthorised access to DMZ servers
2. Poorly secured remote access points

3. Excessive trust relationships in a Domain provide hackers with unauthorised access to sensitive information
4. User or test accounts with weak passwords and excessive privileges
5. Unpatched, outdated software, default configurations
6. Lack of accepted and well defined security policies, procedures and guidelines
7. File and directory access controls
8. Unauthorised services and programs on hosts
9. Weak passwords on workstations
10. Misconfigured Internet server applications and services
11. Misconfigured or poorly updated firewall
12. Running unnecessary services like NetBIOS can compromise network
13. Information leakage can provide attacker with OS type, versions, zone transfers, running services, etc.
14. Inadequate monitoring and detection capabilities at all levels

These are illustrated in figure 1, which shows a typical configuration of a network.




Figure 1: Typical topology of a network with possible weaknesses indicated by numbers.

The steps used to exploit a system follow the general methodology of first conducting an Active Reconnaissance of the network with the aim of *gaining access* by attacking the operating system or conducting an application level attack, scripts attack or targeting default or misconfigurations on the network. Once inside the attacker will try to *elevate the privileges* of the account to allow them to install a backdoor program that will allow future access. Finally they will *cover their tracks* in the system by cleaning

up log files to remove evidence of their presence. If an attacker fails to gain access to the system, they may initiate a Denial of Service attack to prevent anyone accessing the system at all.

Why we need to protect the Network

Networks are pathways to computers and people use computers to store stuff like personal identity, bank, credit card details and purchases made on the Internet. Computers store information ubiquitously on people and companies operate their businesses with and on computers. Two recent reports have highlighted concerns with the way information is treated on computers and the new threats to businesses.

Researchers at RITS Information Security performed a study in how the Irish population dispose of their computers [4] and during this study analysed the contents of recycled hard disks bought openly on the market. The RITS survey found the following:

Organisation Identifiable: In the sample, 33% of the disks originated from the corporate sector ranging from large financial institutions, marketing consultancies, auctioneers, utility organizations, legal solicitors and mobile communication companies. Information included customer's names and addresses, invoices, financial records of past jobs, emails, organisation charts and other relevant documents relating to the organisation.

Personal Information: 62% of the disks were identified as personal computers or home user computers and from half of these could identify their previous owner. This included names, addresses, phone numbers, date of birth and in some cases even bank records and PPS numbers. 10% of the disks contained PPS numbers.

Financial: 24% of the disks contained credit card information. Alarmingly one of these disks contained a spreadsheet of at least 300 credit card details along with expiry numbers, names and addresses.

Passwords: 48% of the disks contained passwords. These ranged from passwords to online sites, email sites, mobile phone sites, etc. These passwords were easily retrieved. No brute forcing of passwords took place.

Illegal Material: 57% of the hard disks contained illegal material.

In the Symantec Internet Security Threat Report 2007 [5] they make the following statements: "The current threat environment is characterised by an increase in *data theft* and *data leakage* and the *creation of malicious code* that targets specific organisations for information that can be used for financial gain" and that "Attackers are now *refining their methods* and *consolidating their assets* to create global networks that support coordinated criminal activity"

This heightened activity in criminal behavior on the Internet is fuelled by the ability to purchase web vulnerability kits and customize them for your exploits.

Be Skeptical: if it is too good to be true, then it usually is

Web vulnerability kits [6] allow an attacker to gain control over client computers when they innocently access web sites hosting the malicious web exploitation kits. The web servers are usually offering free software or games and more often than not, *appear to*

be too good to be true offers. The web server will return malicious malcode as part of the innocent response expected by the client. The newly downloaded malcode will begin a process of installing itself and may access other malicious servers to get more malcode. Now the attacker is in complete control of the client machine and is able to steal personal information from the client and may add the client to a botnet for attacking other computers, possibly in a denial of services attack.




Figure 2: Web Vulnerability Kits have infected many popular websites [6]

IDS – This will protect us, surely?

Intrusion detection systems and firewalls are essentially a detection technology to keep attacks out of your network. They *Detect and Alert* when there are unauthorised access and malicious activities detected in a network. The problems with IDS systems is that it relies on a signature of an attack before the attack can be detected so this can lead to false positives and false negatives, a situation where network traffic is mistakenly blocked or permitted. Also IDS system relies on being able to examine packet headers and payloads in the network traffic, but if encryption is used then it can not be read. Also hackers are constantly using new sophisticated evasion techniques to evade the network security systems.

Honeynets Overview

Motivation

The primary motivation to set up and run a Honeynet is to gather data from attacks and to try and understand the attacks. The main issues are what tools are used, how are they used, by whom and why. What are the tactics and motives of the hackers?

Honeynet Types

There are basically four types of Honeynet deployments; the high interaction Honeynet uses a real network of computers covering a wide variety of operating systems and architectures. The low interaction Honeynets focus on a particular issue like a service attack while virtual Honeynets use a virtual network of computers to simulate a real network. Finally distributed Honeynets are an amalgamation of several Honeynets geographically dispersed to study global attacks.

History of Honeynet Project

The Honeynet Project [7] began in 1999 by several security geeks (as they describe themselves) to investigate the activities of the “bad guys”. Their goals were to learn about tools and techniques and develop new monitoring and counter-attack tools. So much data was gathered that they found it difficult to find time to analyse it all so they created the “*Scan of the Month challenge*” and offered the data openly for anyone to have a go at investigating it. This was so successful that they also created the “*Reverse Challenge*” competition which requires competitors to reverse engineer binary code to analyse malicious applications and code. The Honeynet Project has grown into the Honeynet Research Alliance, a consortium of different academics and professionals that cooperate worldwide in the goals of the Honeynet project.

Brief history of the Honeynets

The Honeynet Project has further developed tools and methods from generation one (Gen-I) to the generation two (Gen-II (2002)) versions and freely distributed this software from their website. The Gen-II tools have significant improvements and together with the benefits of running a Honeywall as a bridge with filtering intelligence give the following features: new tools like SNORT-INLINE, Sebek, rc.firewall, Virtual Honeypots, user interface for management and free bootable CD-ROM images, version 1.1 and mostly recently (June 2007) version 1.2.

The main advances in Honeynets are that as all traffic is suspicious, we have no false positives or false negatives, allow the collection of small data sets, allow the capture of encrypted activity, will work with IPv6, is highly flexible and requires minimum resources to setup and operate.

Honeynet Architecture




Figure 3: One possible layout of the Honeynet

Honeynet Configuration

Important configuration issues are the mode and IP information for the Honeywall, which is set to operate in bridged mode with identification as to which interface is operating as the external and internal bridge interface. The IP address of the Honeynet, the LAN broadcast address and the LAN CIDR address is also needed.

For remote management of the Honeynet, you need to enter the IP address and subnet of HoneyPots, the Gateway address, Hostname, Domain name and DNS (if available), the manager IP address and any restriction on inbound / outbound traffic. Finally the Walleye needs to be activated.

Honeynet Operation

The two essential parts of the Honeywall is data control and data capture. Data Control is being able to provide containment of activity, to monitor inbound/outbound connections, have automatic alerting and the ability to block outward bound activity. Finally all the activities of the Honeywall must be difficult for attackers to detect.

For Data Capture we require monitoring and logging of all activities and data with the challenge to collect as much data as possible without being detected. The Honeywall is layered with the firewall provided by IPTables and the IDS storing full binary data captured of the network traffic using Snort-Inline. When a HoneyPot becomes infected, the attacker's keystrokes are captured and stored on the Honeywall using Sebek.

Before implementing a new Honeywall on a live Internet connection, the following steps should be followed with an offline network to check if Honeywall is working properly;

1. Check if the IPTables logging mechanism is running correctly. You can use a production host to open a telnet connection to the Honeypot. Check to see if the connection recorded. If so, now enable LAN blocking and try connection again.
2. Check if IPTables are limiting outbound connections. Make several HTTP connections to outside world and check the /var/log/messages file on Honeywall. There should be entries with limits noted. Try UDP and ICMP protocols as well.
3. Check Snort-Inline. Use the test.rules file and include it in the configuration file and restart IPS. Now open external telnet session, an HTTP connection and send pings. Check snort-inline alert file for entries and note the dropping of packets and check snort in IDS mode (read data with tcpdump). Finally check for email alerts (if set).

Results from the Honeynet

The ITB Honeynet was setup and configured in June 2007 with Honeybots running standard Windows XP and 2000 client installations. We have found that the Honeybots get compromised very easily and often have to be replaced in the Honeynet so further forensics analysis of the exploit can be made offline.

In this section, we present in detail the data from a typical day of activity on the Honeynet showing the large quantify of data that gets recorded and to give some idea of the types of processes used in an attack.

A summary overview from the Walleye of the 24 hour period of activity is presented in figure 4. It shows the following details:

1. the identity of the Honeywall, date and time of activation and various other localisation information.
2. the top 10 Honeybots
3. the top 10 remote host connections
4. the top 10 source ports and destination ports

This summary information is taken from the recorded pcap packets recorded on the bridge between the Honeynet LAN and the Internet. It is a snapshot of the activity on the Honeynet where no activity should be taking place. The top 10 remote hosts are represented by IPv4 addresses so we can trace them on the Internet. They are unlikely to be the real IP addresses of the attacker because he will have used a compromised computers to do his dirty work for him by activating it remotely and thus avoiding leaving a trail of evidence to him directly.

Honeywall Details for 3238119782							
Sensor ID:	3238119782			Sensor Name:	Honeywall: ITB-AJK1		
Install Date:	Mon Jul 2 10:37:37 2007			Last Update:	Mon Dec 3 20:57:14 2007		
State:	online			Timezone:	0		
Country:	IE			Longitude:			
Latitude:				Network Type:	com		
Notes:							
Activity Report							
Top 10 Honeypots				Top 10 Remote Hosts			
Flags	Host	Connections	IDS events	Host	Connections	IDS events	
	193.1.201.99	345	0	218.15.222.251	3	3	
	193.1.201.102	16	0	202.101.235.100	3	3	
				61.153.194.138	3	3	
				202.99.11.99	3	3	
				218.106.91.25	3	3	
				63.199.210.165	4	1	
				221.209.110.7	8	0	
				193.194.85.74	6	0	
				193.213.34.40	4	0	
				221.208.208.94	4	0	
Top 10 Source Ports				Top 10 Destination Ports			
Port	Connections	IDS events		Port	Connections	IDS events	
1231	3	3		1434	23	15	
1046	3	3		135	59	1	
1143	3	3		1026	449	0	
1084	3	3		1027	395	0	
4431	3	3		1028	365	0	
31093	2	1		138	205	0	
138	205	0		137	121	0	
137	119	0		0	66	0	
0	66	0		445	24	0	
31074	26	0		80	20	0	

Figure 4: Summary of Honeynet Activity for 24 hours

Taking for example the traffic summary report for the 24 hours from the 2nd December to the 3rd December 2007 for detailed analysis, we find that 8,008 packets were processed with a total of 145 IDS events being recorded. The total inbound and outbound packet count is summarised in table 1.

Connection Type	Count	Packets In	Packets Out	Bytes In	Bytes Out
Inbound	1256	1328	0	539359	0
Outbound	8	16	16	0	0

Table 1: Packet Count for 24 hours

The IDS on the Honeywall is configured to limit the number of packets allowed out from the Honeynet from a compromised Honeypot computer. This prevents the compromised Honeypots from engaging in attacks on other computers while still allowing us to examine the attack process in action.

Remote IP	Packets	Bytes	Conns
207.145.74.21	16	0	8
218.10.137.139	7	3199	7
24.64.24.51	6	2904	6
221.208.208.94	6	2742	6
82.71.9.231	10	0	5
221.209.110.50	5	2285	5
62.193.242.99	6	4	4
74.86.42.113	5	0	4
87.67.249.225	4	0	4
220.104.255.79	8	0	4

Table 2: Top 10 Remote IPs:

The source IP addresses were recorded from the packets and these are summarised in table 3. We can trace the origin of these IP addresses using *whois* utility on the Internet. This shows that China and America are the most frequent sources of attacks.

Remote IP	Country
207.145.74.21	United States of America
218.10.137.139	China
24.64.24.51	Canada
221.208.208.94	China
82.71.9.231	United Kingdom
221.209.110.50	China
62.193.242.99	Netherlands
74.86.42.113	United States of America
87.67.249.225	Belgium
220.104.255.79	Japan

Table 3: Countries of origin

The total number of ports scanned (destination ports) was 27 with details of the top ten scanned ports given in table 4 while in table 5 we can see the corresponding applications associated with the ports. Typical attacks are on NetBIOS ports and HTTP ports as well as ping sweeps and MS-SQL attacks. Table 7 and figure 5 show the complete range and frequency of activity on each of the ports. UDP ports of 1026, 1027 and 1028 contain the highest quantity of packets/bytes and connections. The port of 1026 is used by the calendar access protocol and one can suppose that the attacker is trying to use some exploit in applications that use CAP to gain access to the Honeypot.

Port	Packets	Bytes	Conns
udp/1026	418	196110	418
udp/1027	361	174164	361
udp/1028	340	164560	340
icmp/0	101	4197	54
tcp/135	65	9	53
tcp/139	29	0	14
udp/1434	16	4516	14
tcp/22	16	0	11
tcp/445	16	0	8
tcp/80	9	0	4

Table 4: Top 10 Scanned Ports:

Count	SID	Alert Description
4	2	(spp_stream4) possible EVASIVE RST detection
99	384	ICMP PING
2	483	ICMP PING CyberKit 2.2 Windows
2	525	BAD-TRAFFIC udp port 0 traffic
12	2003	MS-SQL Worm propagation attempt
12	2004	MS-SQL Worm propagation attempt OUTBOUND
2	2049	MS-SQL ping attempt
12	2050	MS-SQL version overflow attempt

Table 5: All Snort Alerts

Suspicious Connections:

There is some evidence that the Honeypot has been compromised and has launched attacks on the other computers in the LAN. This is indicated by suspicious activities taking place on the Honeypot even though no users or application services are using the Honeypot. Some activity can be identified as belonging to operating system services that broadcast packets of notification and this can be easily identified and discarded. Other types of activity will require further forensic analysis. Table 5 lists the serious activities on the Honeywall that have been detected by the intrusion detection system, Snort. This list is typical of reconnaissance and footprinting activity performed by hackers. The MS-SQL activity is also typical of a hacker trying to gain access to the system where they would try to download Malware or hacking tools to escalate privileges to account access and take over control of the computer.

Analysis of data for one week (2nd December to 9th December 2007)

Details of the ports used in the Honeynet attacks.

Port	Packets	Bytes	Connect	Shirt	Pocket	netTunes.	Shirt	Pocket
0	602	13,425	334	launchTunes.				
21	24	0	13	FTP, File Transfer Protocol, control.				
22	90	0	64	SSH.				
23	17	0	7	Telnet.				
25	27	0	13	SMTP, Simple Mail Transfer Protocol.				
32	4	0	2					
53	5	0	5	DNS, Domain Name System.				
80	28	0	17	HTTP, HyperText Transfer Protocol.				
110	6	0	2	POP, Post Office Protocol, version 3.				
135	466	15	326	DCE endpoint resolution				
137	19	200	6	NETBIOS Name Service.				
139	140	0	65	NETBIOS Session Service.				
143	3	0	3	IMAP, Interactive Mail Access Protocol.				
443	23	0	14	HTTPS, HTTP over SSL/TLS.				
445	141	0	85	Microsoft-DS.				
1000	6	0	4	cadlock2				
1026	2,093	1,095,178	1,264	CAP, Calendar Access Protocol.				
1027	2,120	923,663	1,156	ExoSee.				
1028	1,965	860,552	1,074					
1070	4	0	2	AT+C License Manager.				
1080	0	3	0	Millicent Client Proxy.				
1433	46	0	25	Microsoft-SQL-Server.				
1434	64	13,164	37	Microsoft-SQL-Monitor.				
2967	8	0	5	Symantec System Center agent.				
3306	4	0	2	MySQL.				
3389	0	2	0	MS WBT Server.				
4899	5	0	5	RAdmin Port.				
5168	7	0	7	SCTE30 Connection.				
5900	25	0	13	VNC Server.				
6101	2	0	1	SynchroNet-rtc.				
8080	16	0	6	HTTP alternate.				
8999	4	0	2	Brodos Crypto Trade Protocol.				
10000	17	0	15	NDMP, Network Data Management Protocol.				

Table 6: List of ports used for attacking Honeypot




Figure 5: Plot shows the frequency of activity on each of the port numbers

Conclusions and Future work

We have shown how easy it is to setup a Honeynet using the Honeynet Project image software and we have collected data that has demonstrated how frequently a computer is attacked and through the suspicious connections we have shown that the basic operating system became infected and compromised on the first day of being connected.

Under Irish law and International law, if your computer is used to attack another computer and gain unauthorised access, then you are responsible and libel for prosecution. So it is important that we are aware of the dangers of being connected to the Internet and how these attacks are conducted and how to safeguard our computers from being compromised.

When we compare our finding with the experiences of other Honeynets, for example the HEAnet Honeynet, we see much of the same activity on the detection level. We are currently working on the forensics analysis of the compromised Honeypots where we are looking for answers to the questions like the following:

- Is the attack real?
- Who is committing the attack?
- What is the timeline?
- Identify all the malicious traffic involved in the attack for offline analysis
- Is there a pattern to the attack?
- What commands / tools were used?
- Was a rootkit used?
- Was an IRC channel used?
- What exploits were used?
- Was the honeypot comprised and used to initiate attacks?

We hope to have completed this work in a few months time and will follow up this paper with the results.

Acknowledgements

The authors would like to thanks the Informatics Department at ITB for the use of their equipment and research network facilities. We would also like to acknowledge the help received from HEAnet in the duration of this project.

References

- [1] The Honeynet Project 2004, “Know Your Enemy”, Publ. by Addison-Wesley
- [2] Siyan, Karanjit, 1997, “*Inside TCP/IP*”, New Riders Publishing
- [3] <http://www.sans.org>
- [4] Vivienne Mee 2007, “Who is reading the data on your old computer”, Journal of Digital Forensics, Security and Law
<http://www.ritspondera.com/>
- [5] **Symantec Internet Security Threat Report 2007**
<http://www.symantec.com/business/theme.jsp?themeid=threatreport>
- [6] Christian Seifert, 2007 “KYE: Behind the scenes of malicious web servers”, The Honeynet Project, www.honeynet.org
- [7] <http://www.honeynet.org/papers/index.html>