Caught in a Honeypot – The Analysis

By

Jay Scott

www.jayscott.co.uk

Introduction

As a security researcher I am always looking for projects to take part in so I started a SSH honeypot project. A great number of Linux-based servers are poorly secured and put on public IP address.

Hackers take advantage of this by running scripts looking for poor server username and password combinations which then carry out a brute-forced attack against the servers. For a hacker to gain access to one of my honeypots they would have to brute-force the root account or the one user account that was set-up.

I researched the data for this paper between the 21st February 2011 and 1st September 2011 using the honeypot software Kippo[1] running on 4 separate Linux-based systems, 3 of which are set-up on VPS servers hosted within the UK, Germany and Sweden and 1 on a UK home cable connection.

All honeypots connect to a central database server where all the data is stored.

Since the goal wasn’t to test strong passwords and the quicker the hackers got access the better I set the root password to something trivial (123456) and the user account had the same password as the username (david/david).

For the project I wanted to focus on 3 main goals which are as follows:

- Report all hacks to the relevant people and record any feedback received, if any.
- Gather as much information on the hackers as possible.
- Collect any tools or malware download to a honeypot by the hackers.

You wish can view all the data and watch the captured hacks used in this paper at my project honeypot website http://honeypot.jayscott.co.uk.

The Hackers

In total there was over a **quarter of a million** (279,322) username and password attempts. The total amount of attacks from **unique** IP address was 2373 and from these unique attacks 751 were able to hack the honeypot and gain root or user access.

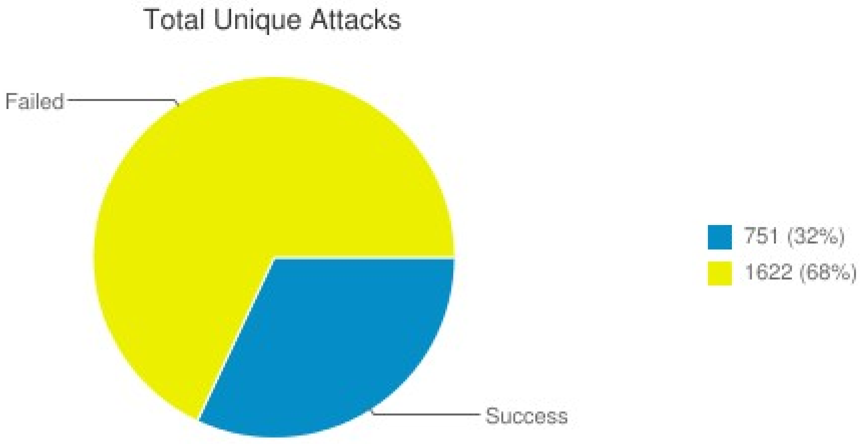

Interestingly from the 751 successful attacks only 329 hackers actually logged in and issued any sort of command. I think the reason for this is due to the hackers launching attacks from other compromised hosts and then not going back to the compromised either because the host had been shut-down or any sort of root-kit/user accounts being removed. Another reason for the hackers not logging into the honeypot could be due to the hackers gathering a list of compromised SSH servers for sale or later use.

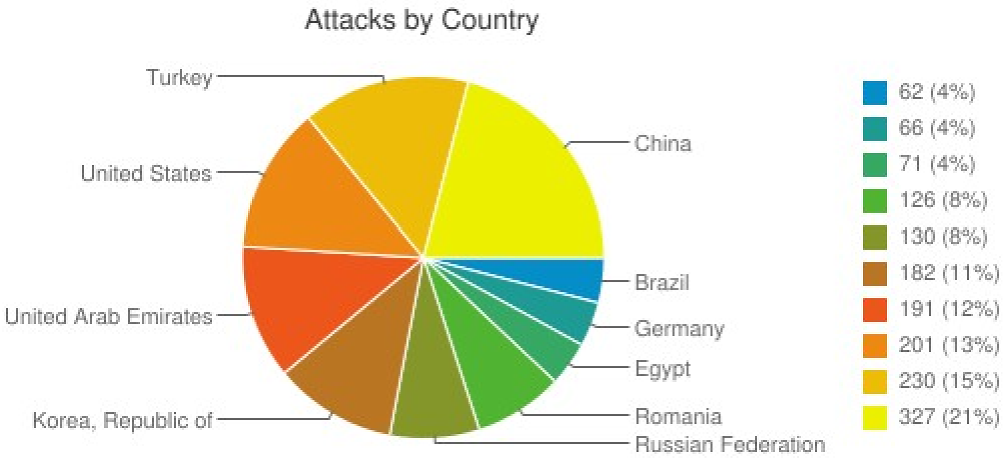

When looking at the top 10 countries where the originating attacks are coming from we can see that China (21%) is the greatest with the nearest country to them being Turkey (15%). Turkey being second on the list was a surprise but recently 32 hackers were arrested in Turkey suspected of DDOS attacks[2].

Malware & Hacker Tools Captured

The honeypots gathered roughly 130 possible malware and hacking tools, unfortunately I was unable to analysis every files content due to time constraints.

However most of time it was easy to see the intended use of the files by watching the commands issued by the hacker. Some of the downloaded files were from Microsoft service packs or 100MB.bin files which I assume are used to test if the server could be used in a DDOS attack or even as a warez dump. The only one file that I could see was different turned out to be a paypal phishing scam.

Most of the tools download would be shell scripts which would execute a SSH scanner and SSH brute force tool. A lot of the scripts were practically the same with only the made by headers changed or simple cosmetic changes to the scripts output. On one occasion a tarred file contained a list of over 3000 Remote Desktop Protocol (RDP) username and passwords along with the IP address, I reported all these IP address.

Reported IP Results

The main goal that I wanted to do for this project is to report all offending IPs to the proper netblock owner which will hopefully stop any attacks originating from that IP. I created a script that would send emails every 24 hours along with the relevant section of the log file to the IP’s owner if they met these rules:

1. Five or more login attempts against a honeypot.
2. A successful login to a honeypot.
3. Report to abuse email[3] on the WHOIS record if found else any email found.

In total 658 unique netblock owners were emailed in trying to stop the hackers. It must be noted that huge amount (40%) of these emails were bounced back with mail configuration errors, unknown user or mailbox full errors.

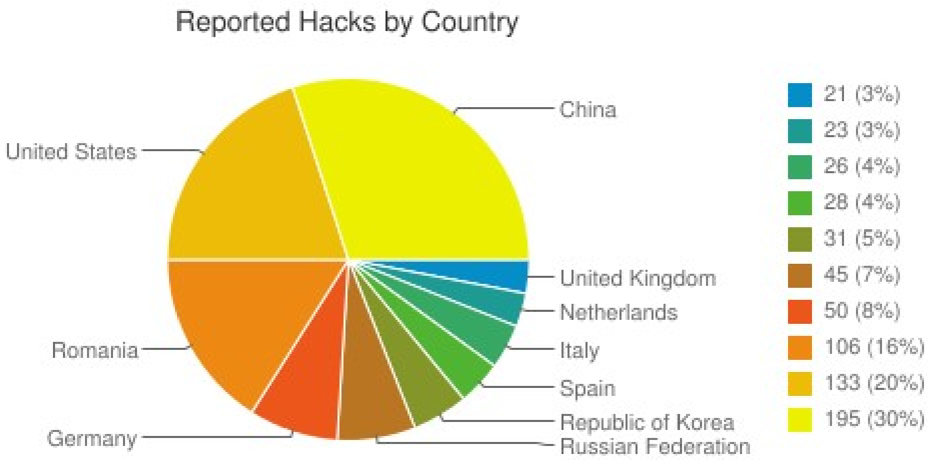

From the 658 netblock owners contacted only 47 replied to the email with any information. Some of the replies would be details on how they are passing the log onto the clients concerned or they have suspended the account while an investigation is being carried out.

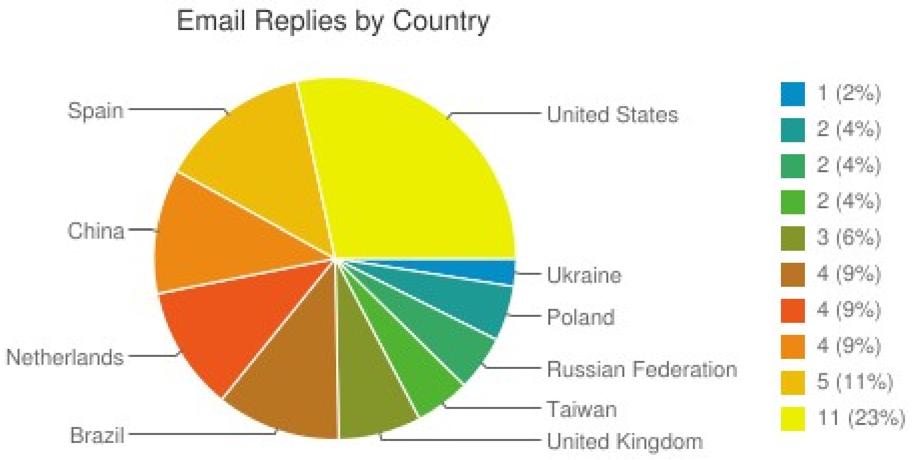

Disappointingly 192 WHOIS records had no email address **at all** so they could not be contacted. The worse offending country for not having WHOIS contact information was Romania (68) very closely followed by America (67).

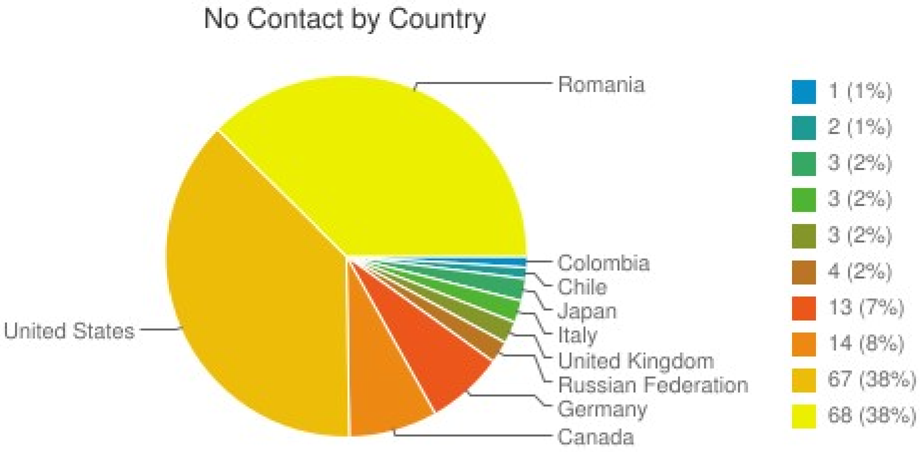

Conclusion

Regarding the actual hack of the honeypot all the hackers seemed focused on of the following:

- Install the SSH scanner used to hack the honeypot to gather more hosts.
- Run a IRC DDOS service for later use.
- Just gathering a list of all hacked SSH servers

Out of all 329 unique hackers zero actually looked around the honeypot for any interesting files or at any important files such as SSH keys. When watching back the hacks you can see that most of the time they seem to be copy and pasting commands into the SSH season, just churning through servers to use later on.

China seems to be where most attacks are coming from but I think this ties in with what I found about reporting IP address. China without a doubt had the highest amount of bounce-back emails this maybe explains why attacks keep coming from the same IP address, the netblock owners just simply don’t know that their servers are compromised.

The most worrying findings from the project for me is the lack of response and the amount of bounce back emails received. 192 out of 946 WHOIS lookups had no email address at all and most of the email address's listed were bounced back as not delivered.

Unfortunately servers do get hacked and if yours does damage limitation is key. However, if there is no way for people to contact you about it and your unaware what’s happening on your network then there is a serious problem.

I am sure people reading this paper know this already but please try to do the following at a minimum to secure your SSH servers and create the best damage limitation if a server does get hacked:

- Keep you WHOIS abuse records up to date.
- Read the abuse email account.
- Don’t set the password the same as the username.
- Don’t use a dictionary based password.
- Do use IPTABLES to block SSH access to only necessary IP address.
- Do change the default port from 22 if you can.
- Do block IP address after X amount of invalid password tries.

If you do the above you have great increased the chances of your SSH server being brute forced.

You can find all the data used on this paper as well as watching live honeypot hacks as they happen at my website: http://honeypot.jayscott.co.uk.

If you which to contact me please email jay@jayscott.co.uk.

References

1. http://code.google.com/p/kippo/
2. http://www.thehackernews.com/2011/06/turkey-police-arrests-32-anonymous.html
3. http://www.ietf.org/rfc/rfc2142.txt