Your Botnet is My Botnet: Analysis of a Botnet Takeover

Malicious code (or malware) has become one of the most pressing security problems on the Internet.Bots are a type of malware that is written with the intent of taking over a large number of hosts on the Internet. Once infected with a bot, the victim host will join a botnet, which is a network of compromised machines that are under the control of a maliciousentity, typically referred to as the botmaster.

There are three main ways to study botnets. First is by performing passive analysis of secondary effects that are caused by the activity of compromised machines such as spam e-mails or DNS blacklist queries. Another solution is by infiltration where someone gets a copy of the botnet, often through honeypots, and then observe the traffic inside a controlled environment(virtual machine). Finally it is possible to change the mapping of a botnet domain to point to a machine controlled by the defender.

Torpig has been distributed to its victims as part of Mebroot. Mebroot is a rootkit that takes control of a machine by replacing the system's Master Boot Record (MBR). This allows Mebroot to be executed at boot time, before the operating system is loaded, and to remain undetected by most anti-virus tools.

After the injection, Torpig can inspect all the data handled by lots of programs and identify and store credentials for online accounts and stored passwords. The Command and Control(C&C) server can reply with two main ways. Okn response is like a server acknowledge. On the other hand there is the okc response, which will send a new configuration file to the bot to perform "man-in-the-browser" phishing attacks. Communication with the injection server is protected using the standard HTTPS protocol. However, since Torpig does not check the validity of the server's certificate and blindly accepts any self-signed certificate. That was a weakness that enabled the researchers to take control of the botnet.

Torpig uses the domain flux to identify and communicate with their C&C servers. It's a list of domain names like "Ren-dezvous points" that may be used by the botmasters to control their bots. In order to achieve that, the botmaster needs to control at least one of the domains that will be contacted by the bots. Also use mechanisms to prevent other groups from seizing domains that will be contacted by bots before the domains under their control. That means that each bot generates a list of domains each week using an algorithm. In practice, the Torpig controllers registered the weekly domain. However, they did not register all the weekly domains in advance, which was a critical factor in enabling the hijacking.

The Sinkholing Preparation was done by purchasing service from two different hosting providers that are well-known to be unresponsive to abuse complaints. During the ten days that researchers controlled the botnet, they collected over 8.7GB of Apache log files and 69GB of pcap data. Torpig obtained the credentials of 8,310 accounts at 410 different institutions.

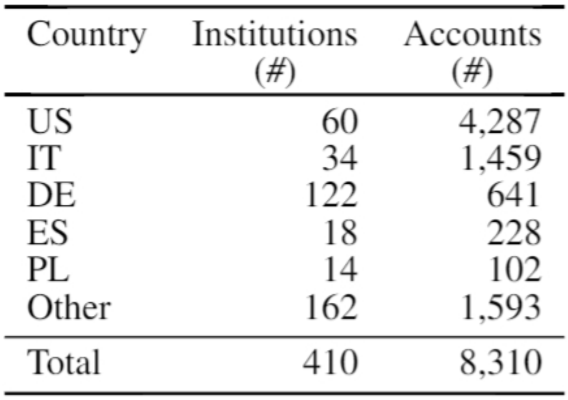

Country	Institutions (#)	Accounts (#)
US	60	4,287
IT	34	1,459
DE	122	641
ES	18	228
PL	14	102
Other	162	1,593
Total	410	8,310

Figure 1: Accounts at financial institutions stolen by Torpig.

In order to measure the botnet size you need the botnet’s footprint, which indicates the aggregated total number of machines that have been compromised over time and the botnet’s live population, which denotes the number of compromised hosts that are simultaneously communicating with the C&C server. By counting unique tuples from the Torpig headers consisting of(nid,os,cn,bld,ver), they estimated that the botnet’s footprint for the ten days of our monitoring consisted of 182,800 machines. In contrast, during the same time, 1,247,642 unique IP addresses contacted our server. So counting the number of infected bots by counting the unique IP addresses that connect to the botnet’s C&C server is problematic. However unique bot IDs per hour with the number of unique IP addresses, they are virtually identical. Thus, the number of unique IPs per hour provides a good estimation of the botnet’s live population.

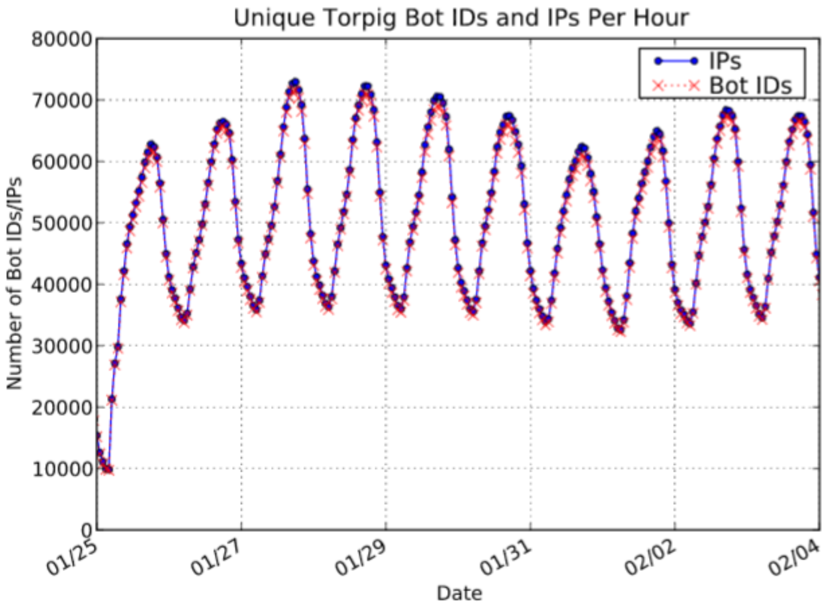

Figure 2: Unique Bot IDs and IP addresses per hour.

In comparison, the number of IPs per day does not accurately reflect the botnet’s live population with a difference of 36.5%. Conversely the median number and average number of IPs per day was 182,058 and 179,866 respectively.

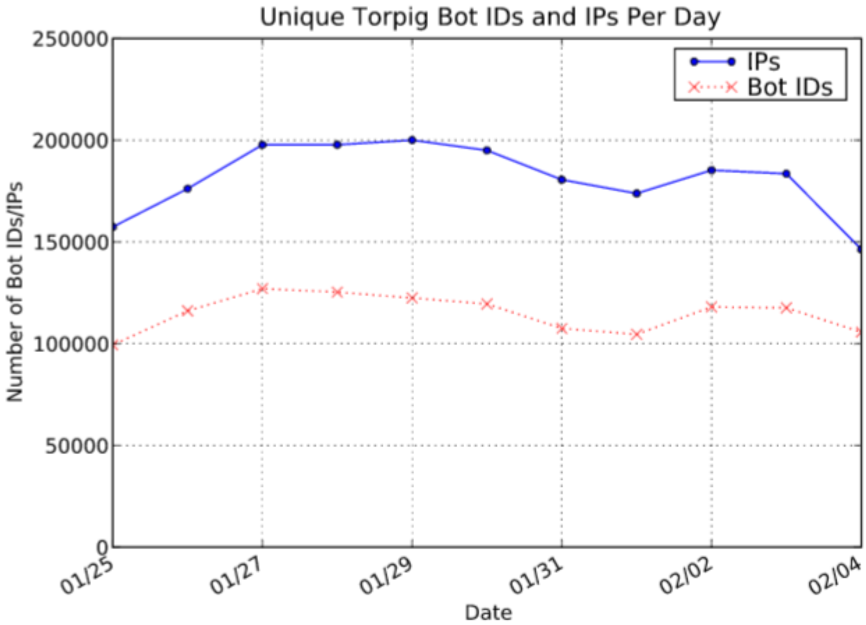

Εικόνα 3: Unique Bot IDs and IP addresses per day.

Dimitriadis Evangelos

CSD
2549