UNCOVERING THE WHOLESALE INDUSTRY OF SOCIAL MEDIA FRAUD: FROM BOTNETS TO BULK RESELLER PANELS

Masarah Paquet-Clouston & Olivier Bilodeau
GoSecure, Canada

{mcpc, obilodeau}@gosecure.ca

ABSTRACT

There is no doubt that there has been an increasing interest in understanding the industry of social media fraud (SMF) – which is the process of creating fake 'likes' and 'follows' on online social networks (OSN) – and its potential deceptive capabilities. This paper explores an undocumented segment of this industry: wholesaling, from botnet supply operations to bulk reselling.

To begin, the paper focuses on a previously unexplored aspect of Linux/Moose, an IoT botnet conducting SMF. Linux/Moose infects devices in order to use them as proxies to relay traffic to social networks. Its architecture includes a whitelist of IP addresses that can push traffic through those proxies – a feature reminiscent of a reseller model. We analyse the traffic fingerprints left by each IP address on the systems we infected and uncover the value of the whitelisted IPs, which is not what we had anticipated. Then, we collect information on bulk reseller panels, the direct working partners of the botnet operators. While analysing their striking similarities, we discover a new key actor in the industry: the software panel seller. We investigate the panels in an attempt to understand how they are connected to main SMF providers like Linux/Moose.

Finally, we map the SMF supply chain, discuss key actors that, if targeted, would disrupt the entire industry and show the likely unequal revenue division in the chain. This is a first review study of the wholesale industry of SMF. It provides key insights for actors willing to curb this illicit activity, from law enforcement agencies to policy makers and cybersecurity professionals.

1. INTRODUCTION

Online social networks (OSN) have become excellent vehicles for marketing, to gain visibility and subsequent influence. They have also become a powerful tool for online manipulation, with erroneous and misleading information being widely distributed on them. Among other fraudulent practices, buying fake visibility through an illicit market for social media fraud (SMF) is a common practice undertaken by many, as reported in journalistic investigations [1, 2] and studies [3–6]. Such practices seem benign at first, but fortunately, law enforcement organizations have started to become aware of the potential for harm they create, such as information manipulation. In January 2018, the New York Attorney General launched an investigation against a social media marketing company, named *Devumi*, that sold SMF to a wide variety of customers [1, 2]. The criminal investigation was set in motion since some of the fake accounts sold appeared to commit identity theft. Yet, there is much more criminal activity behind SMF, and customer-facing companies like *Devumi* are only the tip of the iceberg. The SMF industry involves multiple actors all contributing, to a certain degree, to illicit activities. This explanatory study establishes the link between these actors, from malware authors to final SMF customers, bringing key insights for law enforcement agencies, policy makers and cybersecurity professionals willing to curb the activity of this illicit industry.

To do so, we start by presenting an Internet of Things (IoT) botnet, named Linux/Moose, that conducts SMF on *Instagram*. We study a specific feature of this main SMF provider and attempt to understand how orders are managed in the botnet. We then investigate bulk reseller panels to uncover their middleman role in this industry. Finally, we uncover the potential supply chain behind SMF and its key elements to disrupt the industry and briefly discuss the likely unequal revenue division in the chain.

2. SMF BOTNET: LINUX/MOOSE AND ITS WHITELIST OF IP ADDRESSES

The production of millions of fake accounts to conduct SMF requires automation, which can be done through a botnet – a key actor in the SMF supply chain. Fortunately, in the past two years, we have investigated such a botnet, named Linux/Moose [7]. We begin this section by briefly presenting the botnet, the different steps we undertook to investigate its operations, and the conclusions we have reached so far (with some already published).

Linux/Moose infects embedded *Linux* systems of MIPS or ARM architectures, specifically avoiding x86, such as routers and IoT devices. Just like other well-known IoT botnets, it has a worm-like behaviour, brute forcing Telnet credentials with simple combinations of usernames and passwords. Its main payload is a proxy service that can do SOCKSv4/v5, HTTP, HTTPS, using the infected devices to relay traffic. This allows the botnet operator(s)[1] to hide behind thousands – if not hundreds of thousands – of clean IP addresses[2] [3]. To study the botnet's operations, we built honeypots and infected them with the Linux/Moose malware. We then accessed the content of the encrypted traffic by performing a man-in-the-middle attack using the *mitmproxy* tool [8]. The whole procedure for the honeypots' creation and infection, as well as the man-in-the-middle attack on the encrypted traffic, is detailed in the white paper 'EGO MARKET: When Greed for Fame Benefits Large-Scale Botnets' [3].

Our analysis showed that Linux/Moose's traffic is directed towards various OSNs but mainly to *Instagram*, which represents 86% of the HTTPS requests we studied.[3] On *Instagram*, fake accounts are created and then leveraged to conduct likes and follows on various profiles. Cunningly, to ensure that the fraudulent operations are not caught by the OSN's anti-bot algorithms, human-like behaviours are scripted by the

[1] We do not know how many operator(s) Linux/Moose has.

[2] Reputable IP addresses, likely to be used by legitimate consumers as opposed to data-centre IP addresses.

[3] We use the present tense because we are fairly certain that Linux/Moose is still active.

operator(s) [4]. During the investigation, the profiles followed by the botnet were identified as belonging mainly to members of the entertainment industry (e.g. actors/actresses, singers, newscasters/talk show hosts), small online shops and ordinary people, illustrating that the demand for fake online popularity is vast. For further information, we published a market price analysis of SMF, an evaluation of the botnet's behaviour on *Instagram* and a profile of the potential customers of such fraud in the conference paper: 'Can We Trust Social Media Data? Social Network Manipulation by an IoT Botnet' [4].

The element of interest in this study, which has not yet been published or discussed, is the botnet's distinct use of a whitelist of IP addresses. This whitelist controls who can interact with the bots' proxy and is provided in the regular beacon messages between infected hosts and the command-and-control (C&C) servers. Throughout our monitoring of infected hosts, the whitelist remained the same, containing seven IP addresses. Table 1 shows where each of them is hosted.

Whitelisted IP addresses	Hoster	Country
IP1	CheapWindowsVPS	France
IP2	CheapWindowsVPS	France
IP3	CheapWindowsVPS	France
IP4	ColoCrossing	United States
IP5	Worldstream	Netherlands
IP6	Worldstream	Netherlands
IP7	Worldstream	Netherlands

Table 1: IP hosting providers.

While scanning these whitelisted IP addresses, we found that they were *Windows* servers with their Remote Desktop Protocol (RDP) open, indicating that Linux/Moose could hypothetically be built on a reseller model in which seven distinct servers – and thus potentially seven (or fewer) actors – could log in and use the botnet's infected devices to conduct SMF. This would be a good way to delegate some control to an external actor (partner) without giving that person too much access, thus preventing the botnet from being stolen. We present below the results of our analysis, which does not support the reseller hypothesis, but rather shows a different picture. The analysis involves looking at the traffic fingerprints left by each whitelisted IP address in both the packet capture data (Pcaps) and the *mitmproxy* logs (decrypted HTTPS traffic) on our ten honeypots. Due to the amount of data to process, data visualization[4] was required to find patterns in the traffic. The two subsections below summarize the different findings.

2.1 The similarities in traffic fingerprints of each whitelisted IP address

The results of the five analyses below did not confirm our hypothesis but rather refuted it. They show many similarities in the traffic fingerprints of each whitelisted IP address, forcing us to assume that they are probably controlled by the same actor.

2.1.1 Honeypots used

We looked to see if each honeypot, depending on where it was hosted in the world, was related to a specific IP address in the whitelist. However, as shown in Figure 1, we found no distinctive patterns: each IP sent traffic to almost all honeypots, regardless of where they were hosted in the world.

2.1.2 TLS fingerprints

We then analysed the Transport Layer Security (TLS) fingerprints of each whitelisted IP address to see if different underlying technologies were used by each of them. TLS fingerprints can be found in 'client hello' connections, which is where the connection handshake takes place between the

[4] Using the open source libraries Pandas and Plotly [9].

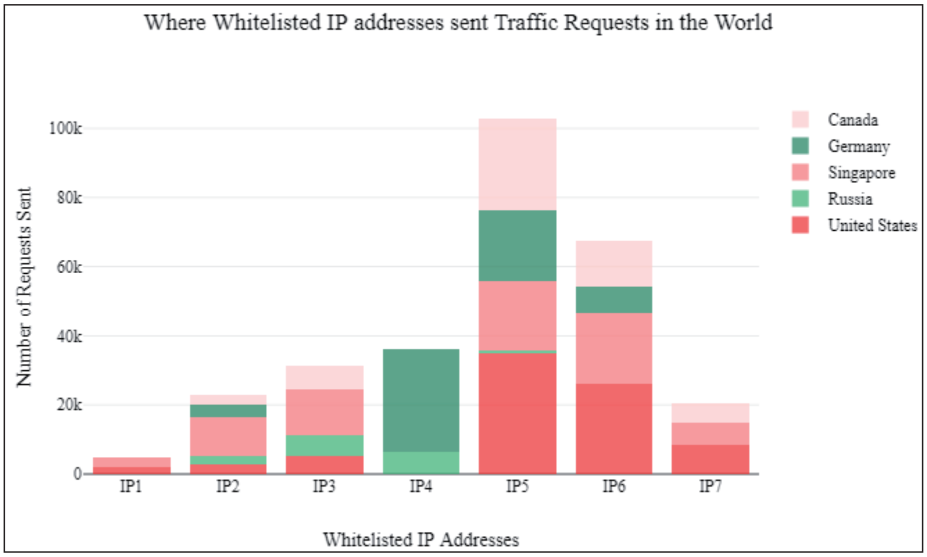

Figure 1: Honeypots.

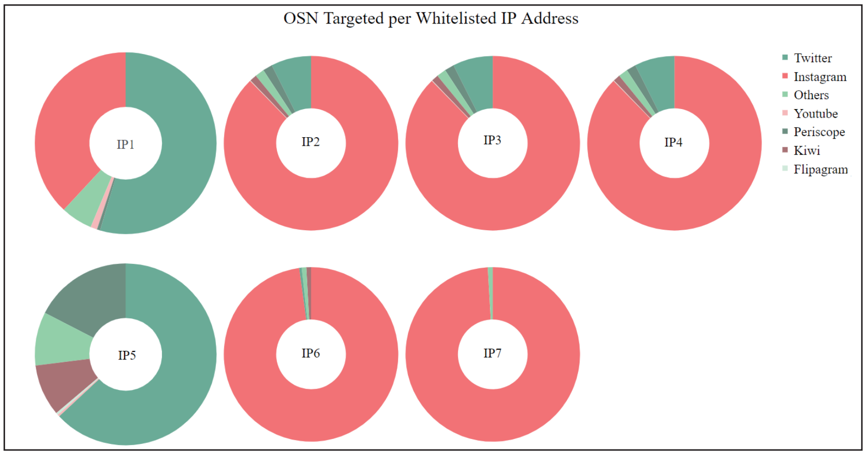

Figure 2: OSN targeted per whitelisted IP address.

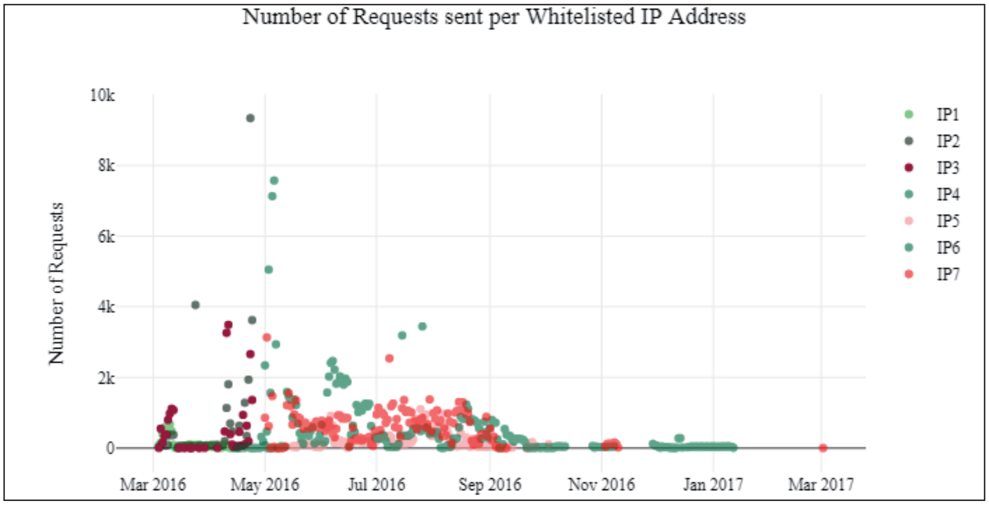

Figure 3: Number of requests in time per whitelisted IP address.

client and the server to find the adequate cryptographic elements that will be used during the connection [10]. In this study, all whitelisted IP addresses used, in most of their connection handshakes, the same TLS fingerprints, except for one. We could not find any distinctive patterns. Furthermore, the TLS fingerprints were not associated with specific libraries or browsers, leading us to a dead end.

2.1.3 Websites targeted

We probed the IP addresses from the whitelist to find out if each was specialized in conducting SMF on a single OSN. Yet, as shown in Figure 2, no specialization was found, as each IP address sent traffic to various OSNs.

2.1.4 User-Agents

We investigated the User-Agent component of each HTTP header, thinking that each actor behind the whitelisted IP addresses could have used different fake User-Agents. We found 3,952 different User-Agents all used by the seven whitelisted IP addresses and varying in terms of browser, phone, library and application types. Most of the User-Agents were used in batches of a few dozen requests, following the scripted patterns for account or friendship/like creations on each OSN. We concluded that the botnet operator(s) are most likely using a tool to randomly spoof the User-Agents in the HTTP headers and that this tool is used consistently throughout all whitelisted IP addresses.

2.1.5 API calls

Finally, we looked at the way the OSN API was used by each whitelisted IP address, focusing on *Instagram* (the main OSN targeted by Linux/Moose). We found that all IP addresses in the whitelist used either the REST or the AJAX API alternatively, with the same sequence of actions imitating humans, as discussed in another study we published [4]. This indicates an identical modus operandi among all whitelisted IP addresses.

2.2 Linux/Moose's use of the whitelist

Fortunately, the analysis of three key variables – timestamps, *Instagram* accounts followed and *Instagram* accounts created –

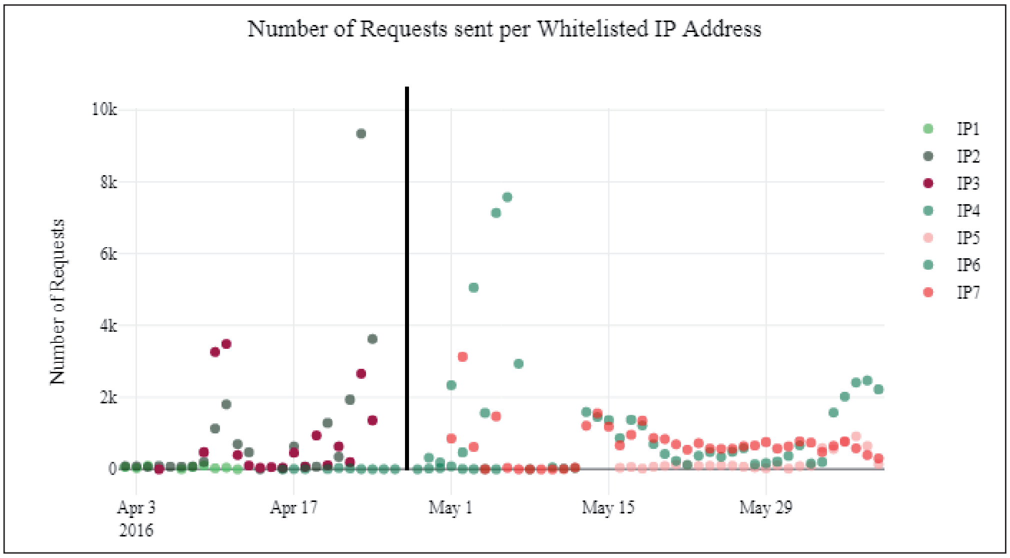

Figure 4: Number of requests in time per whitelisted IP address – April and May 2016.

enlightened us as to the use of the whitelist, which is for fake account management. We also found that three specific IP addresses were active until the end of April 2016, at which point three others took over. This indicated that the operator(s) shifted from one group of active IP whitelisted addresses to another (but left them all in the whitelist).

2.2.1 Timestamps

By graphing the requests over the whole period of study (approximately one year), we found that the seven IP addresses in the whitelist were not used at the same time during that year. Figure 3 shows that some whitelisted IP addresses were active only at the beginning of the study, while others became active only a few months later.

Figure 4 depicts a zoom of Figure 3 for the months of April and May 2016, illustrating that our honeypots received requests from whitelisted IP1, IP2 and IP3 until the end of April 2016, after which requests were received from whitelisted IP5, IP6 and IP7. Only IP4 was used by the botnet before and after the end of April 2016 for a short period of time. This grouping also matches the distribution in hosting services per IP addresses presented in Table 1. From these findings, we could infer two potential events: either a change in resellers had been made or there had been a shift from one group of active whitelisted IP addresses to another (even though all IP addresses remained in the whitelist for the length of the study). The latter turned out to be the most plausible, due to the analysis below.

2.2.2 Instagram accounts followed

We looked at the *Instagram* accounts followed by each IP address in the whitelist and found that many requests aimed at following a specific account were coming from more than one whitelisted IP address. Moreover, we found *Instagram* accounts followed by both the whitelisted IP addresses active before the end of April 2016 and the ones that were active after. This led us to conclude that the findings in the timestamp analysis most likely showed a change in infrastructure from one set of whitelisted IP addresses to another, rather than a change in SMF resellers.

2.2.3 Instagram accounts created

Finally, we investigated the fake accounts created and leveraged to conduct SMF and found that they did not overlap among each whitelisted IP address. Indeed, each IP address in the whitelist had its own unique set of fake accounts. Furthermore, a given fake account was consistently used through the same infected host. This means that, from the point of view of *Instagram*, a fake user always connected to its network from the same IP address – a usage pattern a lot more credible than bouncing around several IP addresses in a botnet.

As mentioned earlier, the whitelisted IP addresses are running *Windows Server* and exposing RDP for remote management. Also, in the case of *Instagram*, although each whitelisted IP address manages a set of fake accounts, there is an overlap in the accounts being followed by the botnet (the potential customers of SMF). This leads us to think that a rather sophisticated proxy-aware *Instagram* fat-client is used to manage the fake accounts and the flows of interaction with *Instagram*. Confirming the exact fat-client that is used by this actor will, however, require further research.

The above analysis allowed us to understand the use of the IP addresses in the whitelist and better clarify the botnet's production process. We now look at other actors involved in the SMF supply chain: reseller panels. These actors were discovered in a previous piece of research [4] that aimed at profiling the demand for SMF by investigating *Instagram* SMF customer profiles found in Linux/Moose's traffic. Some of the *Instagram* profiles followed by the botnet advertised SMF reseller panels, such as http://cheapbulksocial.com/, illustrating that a link exists between these two actors.

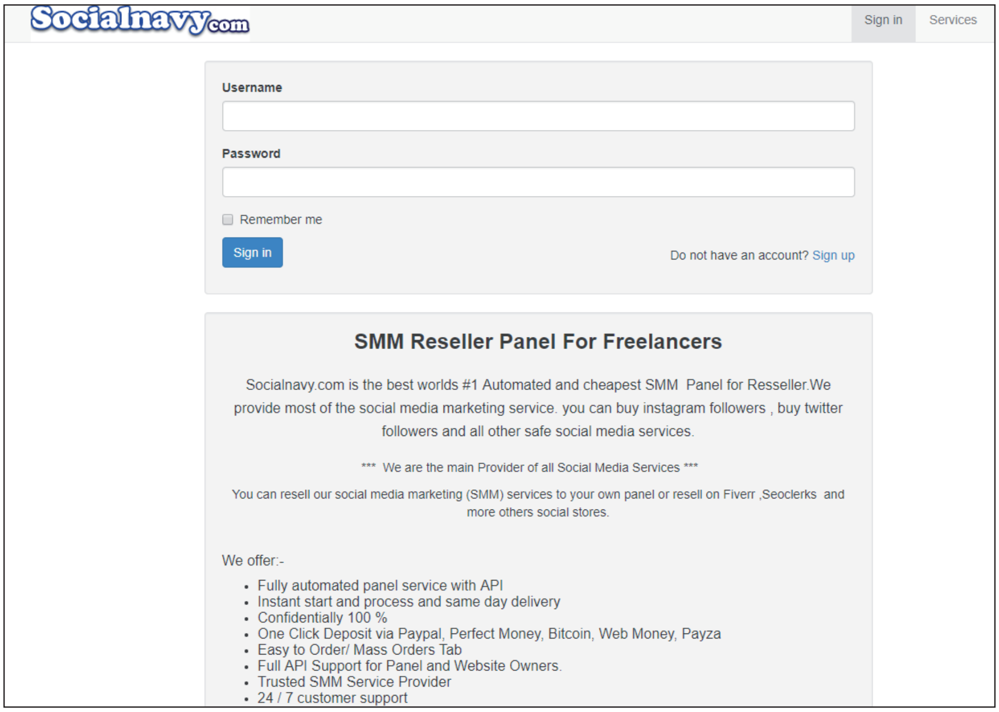

SMM Reseller Panel For Freelancers

Socialnavy.com is the best worlds #1 Automated and cheapest SMM Panel for Resseller.We provide most of the social media marketing service. you can buy instagram followers , buy twitter followers and all other safe social media services.

*** We are the main Provider of all Social Media Services ***

You can resell our social media marketing (SMM) services to your own panel or resell on Fiverr ,Seoclerks and more others social stores.

We offer:-

- Fully automated panel service with API
- Instant start and process and same day delivery
- Confidentially 100 %
- One Click Deposit via Paypal, Perfect Money, Bitcoin, Web Money, Payza
- Easy to Order/ Mass Orders Tab
- Full API Support for Panel and Website Owners.
- Trusted SMM Service Provider
- 24 / 7 customer support

Figure 5: Example of SMF panel.

3. INVESTIGATING RESELLER PANELS

This section presents our brief analysis of reseller panels, which can be considered as the middle-man in the supply chain, connecting customer-facing sellers, like the social media marketing company *Devumi* [11], to main SMF providers, like Linux/Moose. We start by presenting the data collection and the different analytical steps taken to evaluate the striking similarities in reseller panels that brought us to uncover a key actor in the supply chain: software panel sellers. We then investigate software panels in an attempt to understand how they can be linked to the whitelisted IP addresses.

3.1 Data collection

To evaluate these actors, we started by collecting information, searching web pages containing keywords about SMF wholesaling, such as ‘smm bulk’, ‘smm panel’ and ‘smm reseller’ (‘smm’ for social media marketing). We also translated the keywords in Arabic, Turkish, Spanish and Russian to enhance the number of panels found. We discovered a total of 735 web pages that could be related to SMF wholesaling. We went through each of them manually to determine whether the web page was an actual reseller panel or something else. Out of these 735 web pages, 343 were up and represented what we were seeking, as shown in Figure 5.

3.2 Similarities in reseller panels: software panel sellers

Many of the 343 web pages looked similar, raising our first assumption that a small number of actors could be behind SMF reselling. To test this hypothesis, we gathered WHOIS data, certificate information, a fingerprint of the web application (framework, programming language, web container/server), IP addresses and HTML content for each of these 343 web pages and attempted to find similarities in the panels by clustering the data through various methods.

The clustering did not yield much in the way of results, except for one large cluster. We found that 226 web pages (66% of the dataset) were coded in PHP, used similar combinations of client-side JavaScript libraries (including jQuery, Moment.js, Underscore.js, Twitter Bootstrap and Font Awesome) and were hosted on the same IP address belonging to *OVH*. This *OVH* IP address was running an *Ngnix* server on which, according to a query made with the *RiskIQ PassiveTotal Community* platform [12] on 17 July 2018, 977 other domains were hosted. Just glancing at these 977 domains revealed that most of them were related to SMF bulk reselling due to keywords in them, such as ‘smmpanel’, ‘cheapsmmservices’ or ‘bulkfollows’. While visiting them, we found that they were indeed reseller panels with similar login pages, but customized differently (branding, bootstrap themes, pictures, etc.) and showing some variances in prices.

Further investigation of the domains led us to one that advertised the sale of a software panel. The service offered an ‘All-in-one solution for reselling or providing SMM services’, including web hosting and panel maintenance. To use the service, one would only need to own a domain and pay a monthly price based on the number of orders made, ranging from 50 USD up to 200 USD per month.

Moreover, the panel demo available on the website was similar to most of the domains we visited that were hosted on the *OVH*

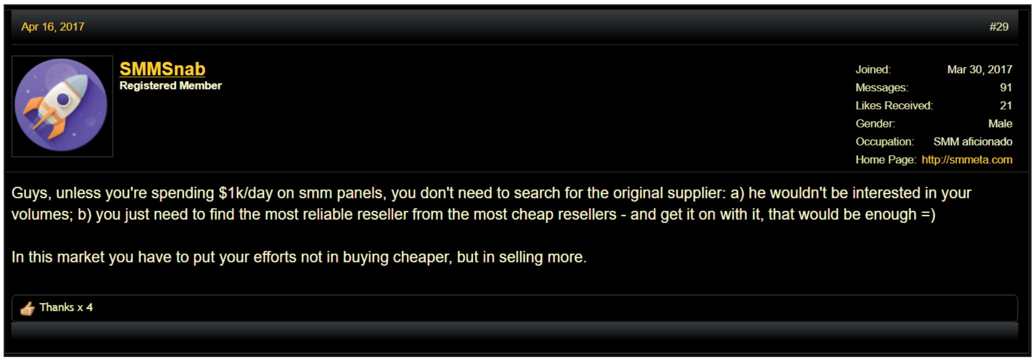

Apr 16, 2017 #29

SMMSnab
Registered Member

Joined: Mar 30, 2017
Messages: 91
Likes Received: 21
Gender: Male
Occupation: SMM aficionado
Home Page: http://smmeta.com

Guys, unless you're spending $1k/day on smm panels, you don't need to search for the original supplier: a) he wouldn't be interested in your volumes; b) you just need to find the most reliable reseller from the most cheap resellers - and get it on with it, that would be enough =)

In this market you have to put your efforts not in buying cheaper, but in selling more.

Thanks x 4

Figure 6: Comment from the thread 'Who Is The Main SMM Panel Services Provider?' on Black Hat World [17].

IP address. This *OVH* IP address could represent the activity of a software panel seller and the multiple clients who have bought the service. To confirm this, we correlated the reCAPTCHA SiteKey found in the signup page of the software panel seller with the ones found in the signup pages of the active domains hosted on the *OVH* IP address. On 17 July 2018, 486 websites were up and were accepting signups (out of the 977 domains hosted on the *OVH* IP address), and from them, 99% used the same reCAPTCHA SiteKey as the one found in the web page of the software panel seller. This allowed us to confirm that the cluster found in the data represented a software panel seller and its related clients.

While investigating the 117 web pages from our first dataset that were not hosted on the *OVH* IP address, we discovered that the cluster found above was not an outlier: there are several software panel sellers available online and they have many clients who all have similar panels. In total, through a similar analysis, we found approximately 1,500 reseller panels and five software panel sellers, three of them offering hosting. For research reproducibility, we published the list of reseller panels and software panel sellers online [13].

Finally, our earlier assumption that only a few actors are behind reseller panels was proven false: there are rather a lot of resellers in the industry, all of them enabled through software panel sellers. The latter are, without question, key actors in the industry, providing panels to those who want to get involved in bulk SMF reselling, but who do not have the technological skills required to build efficient panels.

3.3 From SMF reseller panels to main providers

We tried to find the direct link between SMF reseller panels and main providers, like Linux/Moose. We looked at features offered by three panel software sellers and purchased the service from one of them. We also visited threads on the *Black Hat World* forum related to SMF reselling, to see if we could find any clues as to how resellers find their main providers and how they connect their panels for automated order processing.

We discovered that all software panels provide API access for customers to automatically send orders to them and offer the possibility to complete the received orders manually or use an API to automate the process. The 'automated order processing' API is, however, not linked to a provider. Reseller panel owners need to find their own provider(s). Two software panel sellers mentioned the following feature: 'Automatically Place Orders To Other SMM Reseller Panels'. Resellers can thus connect their panel to other reseller panels and act as an added middle-man in the SMF supply process. Discussions on the infamous *Black Hat World* forum did show that, indeed, many resellers are looking for one or multiple main SMF provider(s) [14–17], but just end up connecting to other, cheaper panels, as suggested in the comment [17] shown in Figure 6.

Unfortunately, apart from the possibility of using another reseller panel's API to fulfil orders, we did not discover how reseller panels fulfilled orders automatically through a main SMF provider. We could not establish the direct link between reseller panels and botnets like Linux/Moose. We know they are related, since some SMF reseller panels were found in the *Instagram* accounts followed by Linux/Moose's traffic. However, the whitelisted IP addresses only had RDP open and thus cannot receive automated API requests from reseller panels. We see two possibilities for them to connect: either the whitelisted IP addresses' aforementioned fat-client fetches the orders by itself or an RDP tunnel is created between the client (such as the *OVH* IP address) and the whitelisted IP address servers. However, further investigation is needed.

4. THE INDUSTRY OF SOCIAL MEDIA FRAUD

Drawing from our findings, we present the potential SMF supply chain, discuss the link between the actors and examine the key ones that could disrupt the industry. We finally provide some insights on the division of revenue in the chain, which seems to be quite unequal.

4.1 The supplying process

Understanding the purpose of Linux/Moose's whitelist, as well as evaluating reseller panels, gave us an interesting overview of the potential SMF supply process, a representation of which is shown in Figure 7.

Figure 7 represents the supply chain from an SMF provider, like the botnet Linux/Moose where the service is produced, up to where the service is sold directly to customers. There are, of course, other potential supply chains. For example, reseller

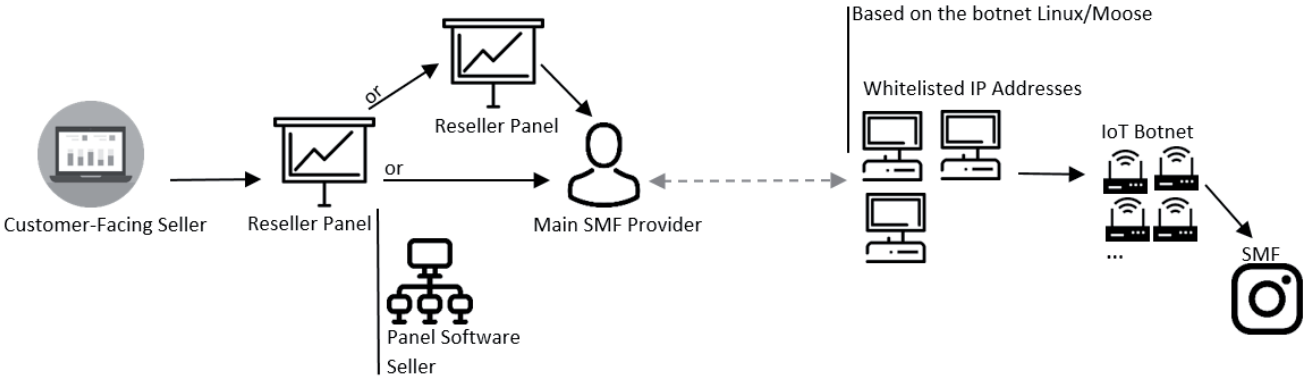

Figure 7: Potential supply chain for the SMF industry.

panels or customer-facing sellers may make their own bot scripts, complete their orders manually, or send their orders to click-farms (large groups of low-paid labourers that conduct SMF manually [18]). Yet, this representation of the supply chain is the most probable one that emanates from this research.

Customer-facing sellers are easily accessible through online searches, as discussed in previous studies of Linux/Moose [3, 4]. They sell SMF services to customers and most likely buy in bulk afterwards from reseller panels. *Devumi*, a customer-facing company currently under criminal investigation [2], was also assumed to buy in bulk: '[...] Devumi does not appear to make its own bots. Instead, the company buys them wholesale – from a thriving global market of fake social media accounts.' [2]

According to this investigation, reseller panels can be linked either to a cheaper panel or to a main provider (through a process not yet uncovered in the case of Linux/Moose). The panel can also be bought from software panel sellers who enable hundreds, if not thousands, of non-technical resellers to get involved in SMF wholesaling. This undoubtedly increases the level of competition among reseller panel owners.

Yet, even though there are more resellers, access to main providers seems to be protected, as software panel sellers do not provide a direct link to them and many resellers on the *Black Hat World* forum ask how main SMF providers can be found. There is no doubt that main SMF providers, like Linux/Moose, avoid any public visibility due to the criminal aspect of their activity. They seem to hide behind multiple actors, staying at the end of the supply chain and dealing only with trusted partners.

4.2 Disrupting the SMF industry

One could decide to disrupt the SMF industry by taking down the main providers in the supply chain – botnets – as they are the ones involved in criminal activities. Yet, such action is a lengthy task that requires technical skills and international cooperation among law enforcement agencies [19]. Other actors in the supply chain are key and may be easier to stop: software panel resellers and reseller panels. These actors enable the link between customer-facing vendors and main SMF providers.

Moreover, by putting an end to the activities of software panel sellers, a large portion of the currently available panels would disappear. This could make the current pyramid scheme crumble and negatively impact many other resellers. Such action would, however, require a thorough criminal investigation of the potential illicit activities behind the panels hosted by software panel sellers, since many options (some of them more licit than others) are available for their associated reseller panels to fulfil orders. Nonetheless, software panel sellers are fundamental to the SMF supply chain: they are the bridge between the non-technical actors (many reseller panels and customer-facing sellers) in the industry and the real cybercriminals behind botnets. Plus, they are not difficult to find as they hide in plain sight.

An easier option would be for the OSNs themselves to go after panel hosters on breach of their 'Terms of Use' instead of going through the criminal route. The timing might be perfect given the political incentives to get rid of online manipulation on OSNs. However, taking panel hosters down might just get everyone into a whack-a-mole scenario where these actors would start hosting their activities on bulletproof hosting providers or compromised servers, reminiscent of what happened with botnet C&C servers.

4.3 Unequal revenue division in the supply chain

Lastly, let's briefly discuss the unequal revenue division found in this supply chain. Indeed, based on a previous study that gathered information on SMF prices of customer-facing sellers [4], the medium price for 1,000 followers on *Instagram* is estimated at 13 USD[5]. On reseller panels, we gathered SMF prices for 58 panels, keeping the lowest and the highest prices for 1,000 *Instagram* followers on each panel[6]. The prices found varied between 0.17 USD (cheap followers) and 15 USD (assumed to be non-drop followers with a specific gender or nationality) for 1,000 *Instagram* followers, with a mean price of 2.74 USD (std=3.30) and a median price of 1 USD. Considering

[5] We take the medium price because the average has a high standard of deviation.

[6] The dataset is available online at [13].

the medium prices of *Instagram* followers for customer-facing sellers and reseller panels, if no other costs are incurred, the former would make a profit of 12 USD for selling 1,000 followers, leading to a net profit margin of 92%. This means that for each dollar invested, customer-facing sellers would receive a net income of 92 cents, a return on investment that is lofty.

On the other hand, such a high profit margin for customer-facing sellers also indicates that reseller panel owners and main SMF providers share a very small portion of what the customer paid for the service. Their profit most likely comes from selling in bulk, as is the case for most producers at the beginning of a supply chain. Yet, for reseller panels, the competition may be quite fierce with many reseller panels available thanks to software panel sellers.

Also, this indicates that main SMF providers earn a very small amount of money for each follower created and sold, as many actors in the supply chain take a large part of the revenue pie. If this business was legal, SMF providers could integrate vertically and own the whole supply chain, allowing them to absorb the profit margin from other actors. Yet, the illicit aspect of their activity forces them to stay in the shadow and earn less per follower produced.

Lastly, software panel sellers offering hosting seem to have an alluring business. For every panel they host, they earn a fee per month, ensuring recurring business from their customers. The fee, at the time of writing, was around 25 USD to 200 USD per month, depending on the number of orders completed on the panel. With these fees, few subscriptions are needed to start making an interesting amount of money. For example, hosting 100 panels at a minimum of 25 USD would give a revenue of 2,500 USD per month. Considering that the *OVH* IP address hosted 977 domains (not all of them active, however), this software panel seller may make an interesting monthly revenue.

5. CONCLUSION

In this paper, we have uncovered the potential SMF supply chain and discussed the operations of multiple actors involved. Further research is, however, required to better understand the dynamics that link each actor in the supply chain. For example, how many reseller panels link their panels to others and how many have direct access to main SMF providers? Are some reseller panels more successful than others, and if so, why? Also, we still have not found how reseller panels are linked to Linux/Moose's whitelisted IP addresses.

Still, this research is part of a two-year-long investigation that attempts to place a botnet's operations in its economic context, understanding the ecosystem in which it evolves. This specific study is a first review of the wholesale industry of SMF and provides key insights for actors willing to curb these fraudulent activities.

ACKNOWLEDGMENTS

We would like to thank Philippe Arteau and Benjamin Mbemba for taking the time to review this paper.

REFERENCES

[1] Harris, R.; Confessore, N.; Dance, G.J.X.; Hansen, M. 2017. Astroturfing, Twitterbots, Amplification – Inside the Online Influence Industry. https://www.thebureauinvestigates.com/stories/2017-12-07/twitterbots.

[2] Harris, R.; Confessore, N.; Dance, G.J.X.; Hansen, M. 2018. The Follower Factory. New York Times (January 2018). https://www.nytimes. com/interactive/2018/01/27/technology/social-media-bots.html.

[3] Paquet-Clouston, M.; Bilodeau, O.; Décary-Hétu, D. 2016. EGO MARKET: When Greed for Fame Benefits Large-Scale Botnets. Technical Report. GoSecure Technical Report.

[4] Paquet-Clouston, M.; Bilodeau, O.; Décary-Hétu, D. 2017. Can we trust Social media data? Social network manipulation by an IoT botnet. In Proceedings of the 8th International Conference on Social Media & Society. ACM, 15.

[5] Stringhini, G.; Wang, G.; Egele, M.; Kruegel, C.; Vigna, G.; Zheng, H.; Zhao, B.Y. Follow the green: growth and dynamics in twitter follower markets. In Proceedings of the 2013 conference on Internet Measurement Conference. ACM, pp.163–176.

[6] Thomas, K.; McCoy, D.; Grier, C.; Kolcz, A.; Paxson, V. Trafficking Fraudulent Accounts: The Role of the Underground Market in Twitter Spam and Abuse. In USENIX Security Symposium. 2013. pp.195–210.

[7] Bilodeau, O.; Dupuy, T. 2015. Dissecting Linux/Moose. http://www.welivesecurity.com/wpcontent/uploads/2015/05/Dissecting-LinuxMoose.pdf.

[8] mitmproxy. https://mitmproxy.org/.

[9] https://pandas.pydata.org/ and Plotly https://plot.ly/.

[10] Brotherston, L. TLS fingerprinting – Smarter Defending & Stealthier Attacking. 2015. https://blog.squarelemon.com/tls-fingerprinting/.

[11] Devumi. https://devumi.com/.

[12] RisqIQ Community. https://community.riskiq.com/.

[13] https://github.com/Masarah/ResellerPanels.

[14] Black Hat World. 2017. Original SMM Panel Provider. https://www.blackhatworld.com/seo/original-smm-panel-provider.986779/.

[15] Black Hat World. 2017. Smm panel supplier (the boss). https://www.blackhatworld.com/seo/smm-panel-supplier-the-boss.993524/.

[16] Black Hat World. 2017. Who can tell me the SMM service provider panel rather than the reseller panel? https://www.blackhatworld.com/seo/who-can-tell-me-the-smm-service-provider-panel-rather-than-the-reseller-panel.980898/.

[17] Black Hat World. 2017. Who Is The Main SMM Panel Services Provider? https://www.blackhatworld.com/seo/who-is-the-main-smm-panel-services-provider.867369/page-2.

[18] Nguyen, C. 2017. What Life for a Bangladeshi Click Farmer Looks Like. https://motherboard.vice.com/en_us/article/9a3az3/ what-life-for-a-bangladeshi-click-farmer-looks-like.

[19] Paquet-Clouston, M.; Décary-Hétu, D.; Bilodeau, O. 2018. Cybercrime is whose responsibility? A case study of an online behaviour system in crime. Global Crime 19, 1 (2018), pp.1–21.