Tracking Botnets on Nation Research and Education Network

Ivan Daniel Burke, Dr Alan Herbert
Council for Scientific and Industrial Research, Pretoria, South Africa
Rhodes University, Grahamstown, South Africa
iburke@csir.co.za
a.herbert@ru.ac.za

Abstract: The South African National Research and Education Network (SA NREN) proves network connectivity and services to all tertiary education networks and research councils within South Africa. The NREN forms part of South Africa's national integrated cyber infrastructure, as such, it is a potential target for cyber-attacks. Due to the large volume of traffic and decentralised nature of the SA NREN, monitoring, reporting and mitigating cyber-attacks is a complex problem. The NREN Cyber Incident Response Team (CSIRT) uses network flow data to identify early indicators of cyber-attacks. In this paper the focus will be on the mechanisms used to identify malicious botnet traffic using network flow analysis.

Keywords: Network flow analysis, NREN, Network Traffic Analysis, botnet detection, cyber threat detection.

1. Introduction

The South African National Research and Education Network (SA NREN) provides the backbone infrastructure for most of South African research and education institutions. This network operates similar to the American, ESnet, or the European SURFnet and CESNET organisations. These NRENs provide services for these research institutions such as network interconnectivity, large data transfer capability, data warehousing and data processing services. In the SA NREN, the backbone is maintained and secured through the collaboration of two institutions: the Tertiary Education and Research Network of South Africa (TENET) which is a service organisation which maintains the services running within the NREN and the South African National Research Network Competency Area (SANReN CA) both these organisations form part of the National Integrated Cyber Infrastructure System (NICIS) organisation.

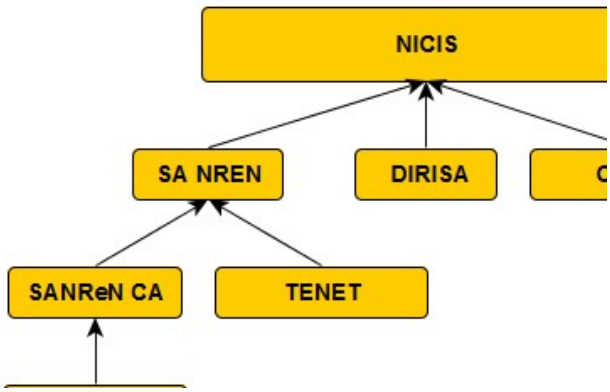

Figure 1: NICIS organogram

Within the SANReN CA a Cyber Security Incident Team (CSIRT) was established in 2016 to address the growing concern of cyber threats against the NREN. Within the CSIRT various tools and mechanisms are used to detect network anomalies and investigate cyber-attacks. The CSIRT aims to provide a proactive cyber incident capability to the SA NREN user base. Thus the CSIRT monitors the NREN for signs of anomalies or abuse before it is reported to the SANReN CSIRT team. As part of this capability the CSIRT seeks to detect: Port scans, Brute force attacks, Denial of Service and Botnet detection. However, this paper focuses the work done within the CSIRT, regarding botnet detection using network flow analytics.

The reason for using network flow analytics is primarily due to the reduced storage size requirement as appose to storing raw packet data. This has the additional benefit of protecting user privacy. Each academic entity within the NREN is entitled to their privacy. Network flow data allows the CSIRT to monitor the traffic flows without violating the user's privacy. The negative trade off of not capturing data payloads is that it makes certain attacks, such as phishing email and drive-by downloads, nearly impossible to detect.

Not all institutions within South Africa have similar budgets to address ICT or cyber-crime related issues. The SANReN CSIRT aims to supplement existing institutional ICT capabilities where needed.
This paper describes the research effort by the SANReN CSIRT to detect Botnets by addressing the following key topics. Firstly the current research methodologies to detect malicious network activity is discussed in section 2. Next, a brief overview of active botnet families, during the observation period is presented in section 3. In section 4 the network flow collection environment is explained. Section 5 showcases the results achieved by applying the research to our data set. In section 6 the paper concludes with our findings and recommendation for future work.

2. Related Work

For the purpose of this paper a botnet is defined as a collection of infected hosts (computers, routers, IoT devices, etc.), known as bots, which are controlled by a bot master via a collection of C2 servers to perform coordinated cyber-attacks. Based on the targets and C2 structures used these botnets can be classified into various botnet families or strains.
Amini, et al. (2014) created a categorisation of techniques used to identify botnets using network flow data. These categories were as follows:

- Correlation: These techniques uses statistical analysis to identify correlations between user access patterns and the number of unique flows. The DISCLOSURE system proposed by Bilge, et al. (2012) is an example of using correlation to detect botnet activity.
- Clustering: This technique aims to cluster together similar traffic patterns into groupings of similar behavioural patterns. This is the technique used by Amini, et al. (2014) to identify botnets.
- IRC channel monitoring: IRC is a well-known communication channel used by botnets to perform Command and Control (C2) activities (Feily, et al., 2009, Amini, et al., 2014). These techniques focus on behavioural fatteners of IRC communications (Mazzariello, 2008). Network flow analysis does not provide new insight into the detection of IRC botnet detection and previous work by other research have covered these techniques in great detail (Mazzariello, 2008).
- Machine learning algorithms: Due to the fact that Machine Learning (ML) algorithms are widely used to identify patterns in large data sets and to perform clustering on seemingly disjoint data sets (Wagner, et al., 2011). ML algorithms was used by Bilge, et al. (2012), Hofstede (2013) and Amini, et al. (2014) to assist with clustering and correlation detection.
- DNS request monitoring: Botnets seek to hide their C2 server infrastructure by using Domain Generating Algorithms (DGA) and Fast-Flux techniques (Grill, et al., 2015, Erquiaga, et al., 2016). This results in an abnormally high number of DNS look-up requests from hosts which are infected by bots.

DISCLOSURE is a botnet detection platform which uses network flow analysis to identify C2 servers within a network (Bilge, et al., 2012). The system uses the random forest machine learning algorithm to classify servers on a network as either benign or malicious. The DISCLOSURE system achieves this classification by first computing several features of the network flow data. The feature extraction algorithms focuses on flow-size, client access patterns and temporal behaviour as input for the classifier.
The flow size features used by the DISCLOSURE system are the unique flow sizes (number of bytes) observed during a period, statistical features such as mean flow size and standard deviation from the mean flow size. The researchers also used autocorrelation to derive additional statistical features based on the flow size (Bilge, et al., 2012).
Bilge, et al. (2012) theorized that benign user interaction with a server results in greater variation in flow sizes than that of a botnet C2 communication. Thus it would be expected that the standard deviation observed by benign servers would be greater that the deviation observed for C2. Furthermore benign servers should have a higher number of unique flow sizes compared to C2 server. These features assisted the classifier in classifying flows as either benign or malicious.
DISCLOSURE used client access patterns to further classify potential C2 servers. Bilge, et al. (2012) monitored regular access patterns by creating a data series wherein the time delay between consecutive server access attempts from each client is captured. Then statistical analysis was performed based on the inter-arrival times of each client. Bilge, et al. (2012) theorized that bots communicate with C2 server in short regular burst in order to avoid detection.
The final feature taken into consideration for the DISCLOSURE classifier is the temporal feature related to the time of the day it is expected for users to be more active. Thus it is uncommon for users to be equally active

during the early hours of the day or late at night compared to the activity observed during the middle of the day.

Amini, et al. (2014) used GNS3 network simulator to create a test environment for their cluster testing. Amini, et al. (2014) infected nodes within the network with the Zeus bot. Amini, et al. (2014) created their clusters by first running the simulated environment for a time period whilst collecting the network flow data. After concluding the experimental run benign, traffic was filtered based on whitelist rules and black list rules, resulting in two disjoint data sets. These data sets were then used to identify C2 flow events and create behaviour rules for identifying C2 communication activity. A critique of the methodology used within Amini, et al. (2014) work is that only one family of botnets was tested (namely Zeus) and the test was conducted on a small test network, less than ten nodes, with a known infection. During our own testing, the results obtained by Amini, et al. (2014) could not be replicated in a larger network were the existence of a botnet was unknown.

Hofstede (2013) used the Exponentially Weighted Moving Average (EWMA), a ML algorithm, to detect network intrusions in real time. The EWMA algorithm reduces the significance of previous measurements based on the amount of time which has passed since the measurement was taken. Thus the weight of a measurement on the average measurement is reduced over time. Hofstede (2013) introduced a supplementary algorithm which they refer to as a seasonal EWMA. The seasonal algorithm takes into account that network traffic usage is influenced by the user's utility of the network at a given time of day. Thus Hofstede (2013) selected a seasonal period of one day and averaged the EWMA results with the measurement taken at the same time the previous day. The seasonal algorithm extension is an alternative to the temporal feature extraction proposed Bilge, et al. (2012), both approaches consider the influence the time of day may have on network utility.

Hofstede (2013) collected data from the CESNET backbone network over a 14-day period in 2012.

3. Active Botnets During the Observation Period

The network attack detection research within the SANReN CSIRT started in November 2017. Based on our current findings a large volume of the attacks are directed at Internet of Things (IoT) devices and specifically Huawei networked devices. This is likely due to Huawei devices being specifically targeted by several botnets (Wijesinghe, et al., 2015, Antonakakis, et al., 2017). According to GlobalStats (2019), Huawei control 25.82% of South Africa's mobile device market share, see Figure 2.

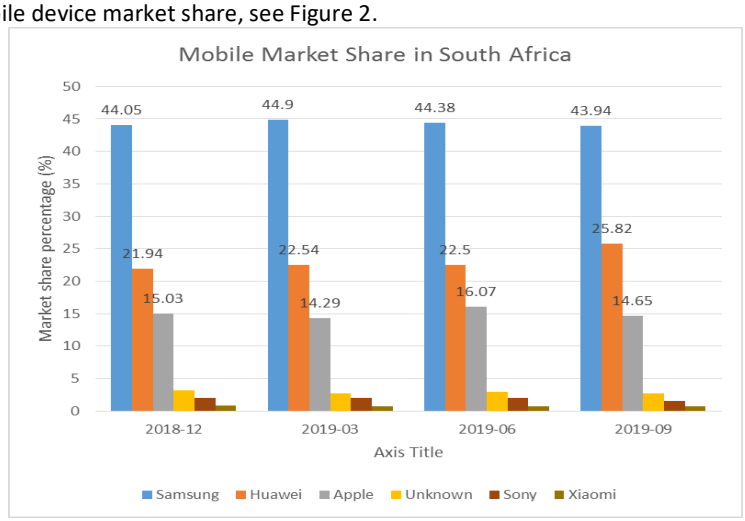

Figure 2: South African mobile market share

Since the start of this research several notable botnet families were observed in the wild. Since the SANREN CSIRT currently only captures IPFIX data we were unable to directly map the cellular operating system to the botnets observed during these attacks. However due to the large market share of Huawei devices and a number of botnets specifically targeting these devices the research conducted within the SANREN CSIRT focused on botnets which are known to target these devices. In this section a brief overview of the main botnet families will be provided.

3.1 Mirai Botnet

The earliest traces of the Mirai botnet was observed in August 2016. The Mirai botnet spreads like a network worm, it primarily targeted Internet of Things (IoT) devices. The most notable attacks performed by the Mirai botnet is the Distributed Denial of Service (DDoS) attack against OVH and Dyn.
The botnet, originally infected new hosts by performing TCP SYN probes on TCP port 23 and TCP 2323 (Telnet). If the target responded to the TCP SYN request, the botnet would enter phase 2 of the infection step by performing a brute force login attack on the service. The Mirai botnet had a list of 62-preconfigired credentials which were known default credentials for vulnerable IoT devices. Once the bot successfully logged into the device a control message would be sent to a Mirai C2 server and a target specific malware payload would be installed on the newly infected device. The newly infected device would then start then start seeking out new devices to infect. Antonakakis, et al., (2017), identified 112 C2 domains and 92 IP addresses related to the Mirai botnet by using active and passive DNS.

3.2 Satori Botnet

Over time Mirai expanded its infection pattern by targeting TCP port 22 (Secure Shell, SSH), 80/8080/443 (web interface, HTTP/S), 7547 (CPE WAN Management Protocol, CWMP), 5555 (Android Debugging Bridge, ADB). The Satori strain, of the Mirai botnet family, exploited CVE-2014-8361 and CVE-2017-17215 via ports 37215 and 52869 to exploit Huawei HG532e home routers without relying on the brute-force attack phase (Anubhav, 2018). In July 2018, it was reported Satori targeted Android devices by sending a malicious payload to the ADB using port 5555(Pour, et al., 2019).

3.3 GoldBrute Botnet

The GoldBrute botnet performs a brute-force attack on servers which have TCP port 3389 (Remote Desktop Protocol, RDP) exposed to the internet (Marinho, 2019). In June 2019, it was reported that the GoldBrute botnet is performing an active attack on over 1.5 million RDP devices. According to Shodan, there are 3.8 million devices (27 589 of which are in South Africa) currently connected on the internet which expose TCP port 3389.
This botnet was controlled via a single C2 server, 104.156.249.231. Bots connect to the server over TCP port 8333. Once the botnet successfully brute-forces a RDP server, the server is instructed to download the botnet payload from the following IP 104.248.167.144. The payload is an 80 megabyte Java application (Marinho, 2019). After the application is downloaded the newly infected server will form part of the botnet and seek out other RDP servers to infect. The GoldBrute C2 server was discovered and taken down in July of 2019.

4. Network flow data collection within SA NREN

The network flow data is captured within the IPFIX data format. IPFIX is a vendor neutral format for storing network flow data, formerly known as NetFlow. The network flow data contains only the meta-data related to a network connection stream, with none of the packet content. At the bare minimum a network flow can be captured as a 6-tuple of information elements: source IP address, destination IP address, source port, destination port, protocol used and time-stamp. IPFIX is based on of NetFlow version 9 which uses 13-tuple to represent a network flow. Each network connection is usually represented by its own flow. These flows are exported to the IPFIX format by an IPFIX export device (usually routers or switches). In order to centralise the analysis process an IPFIX collector is used to fetch all IPFIX logs to a centralised server. Within the SA NREN there are ten primary IPFIX exporters which are used for our analysis. Each exporter is external to the research institutions and each collector collects the data from various institutions. Most of the IPFIX exporters export the flows at a 1:1 ration, however some of the exporters only export a sample of the observed flows. Table 1 shows a list of network flow exporters within the SA NREN as well as their average flows. Each exporter exports well over a billion flows per day (These were the averages observed during the period 1 May 2019 till 31December 2019). Processing such a large volume of data does present a significant challenge to the SANREN CSIRT.

Table 1: Network flow exporters within SA NREN

Device ID	Ratio	Average flows per day
AMS1-IR1	1:1	2.44×10^{10}
CPT1_IR1	1:1	2.66×10^{10}
CPT1_PE2	1:50	6.37×10^{9}

CPT3_PE1	1:50	1.64×10^9
CPT3_PTA1	1:50	2.14×10^9
ISD1_PE1	1:50	1.89×10^9
JNB1_PE1	1:1	8.07×10^9
LDN1-IR1	1:1	1.83×10^{10}
MTZ1_PE1	1:1	1.62×10^{10}
PTA1_PE2	1:1	4.94×10^9

On average 1 terabyte of network flow data is captured each month. Due to the large volume of data to be processed a detection mechanism which is capable of scaling to the needs of the SA NREN is required.
The SANReN CSRIT has an additional constraint in that the network flow data is captured on the SA NREN backbone infrastructure rather than within the specific research institutions themselves. This introduces constrains such an aggregated flows whereby network traffic flow through network aggregation points such as institutional firewall or Network Address Translation (NAT) devices. If flows are aggregated the true source or destination of a network flow may be obfuscated from the network flow logs.
During our initial investigations it was observed that some institutions inject additional packet overhead into their network flow data. This data tends to be related to parity checks, GRE or NAT information. These alterations to the true packet size may influence the detection of malicious traffic if traffic size is used as an indicator.
Another constraint placed on the SANReN CSIRT is that the NREN infrastructure is owned by the constituents and that the CSIRT can only provide advisories to the constituents. The CSIRT has no authority to take down any misbehaving nodes or block connectivity to potentially malicious nodes. In extreme cases TENET can be tasked with limiting network access but this is only as a last resort.

5. Applying research to practise

From the SANReN CSIRT's perspective there are four main activities which are monitored using IPFIX data:

- Port scans detection: Port scans are very prevalent on large scale networks and are not of critical interest to TENET, from a threat perspective. The greater concern is trends or sudden spikes in specific ports being scanned.
- Brute-force login attack detection: Brute-force login attempts fall into a similar category of concern as port scans, in that they are prevalent on the SA NREN. However, brute-force attacks pose a significantly higher risk than port scans if a system is compromised. From TENET's perspective detecting and preventing brute-force attacks aimed at key backbone infrastructure nodes are the highest priority for detection.
- Distributed Denial of Service Attack detection: The IPFIX is used to identify potentially malicious nodes inside the SA NREN which participate in DDoS activities. Due to the meta data capture using network flows volumetric attacks are easily detectable. The SANReN CSIRT is more interested in protocol specific attacks. Specifically Distributed Reflection Denial of Service (DRDoS) attacks which uses SA NREN infrastructure to launch attacks against external services or service providers.
- Botnet detection: From the SANReN CSIRT perspective the team is interested in detecting and monitoring botnet activity at a national level, including the spread of infection within tertiary institutions. TENET is concerned about their infrastructure being targeted by a large scale botnet attack or inadvertently hosting bots on the tertiary infrastructure.

This paper focus specifically on the botnet detection component of the SANREN CSIRT efforts, however the detection efforts rely substantially on the prior work done to detect other threats to the SA NREN.
TENET is responsible for AS2018 (https://ipinfo.io/AS2018), as The SANReN CSIRT needs to monitor several large blocks of IP addresses. In this section several incidences which occurred on these network ranges will be discussed. In order to mask the various institutions identities the IP blocks have been mapped to the private IP address space 10.0.0.0/24.
The SANReN CSIRT does not rely on a single algorithm to detect abusive behaviour or cyber-attacks. Instead, the CSIRT rely on a combination of several algorithms to identify potentially malicious IP addresses. The CSIRT maintains a list of IPs of Interest (IPoI) which are tagged with labels, describing their potential malicious behaviour. The idea of tagging of IPs based on their behaviour was first proposed by Sweeney & Irwin (2017). For example if an IP is observed scanning port 22 (SSH) that IP will be tagged in the database with the tag "Port

scan - 22". IPs within the database can be assigned multiple labels since it is likely that a misbehaving node could potentially perform multiple suspicious activities.
The SANReN CSIRT uses the IPFIX port scan signatures proposed by Grégr (2010) to classify port scan activity based on TCP flag combinations

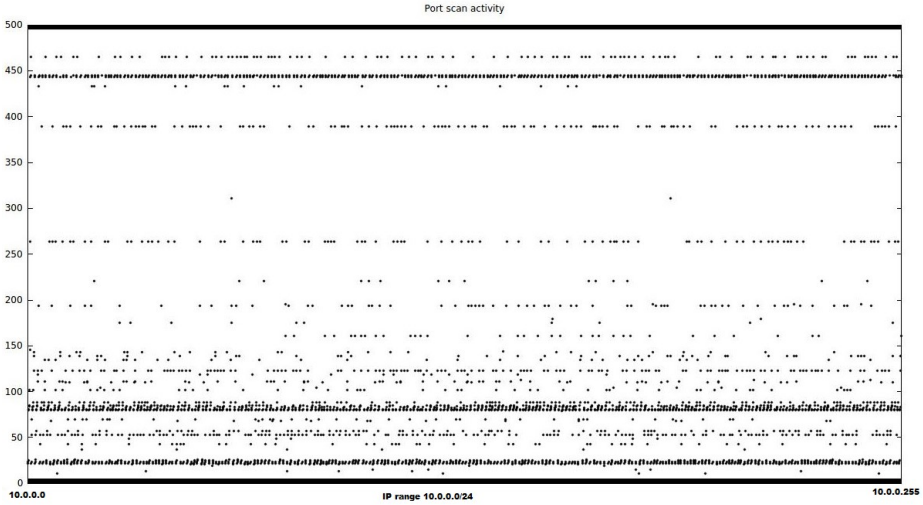

Figure 3: Stealthy port scan activity observed over /24 network range

Figure 3 depicts a scatter diagram of an IPoI performing a SYN scan on a /24 network range within the SA NREN. The y-axis represents the ports being scanned and the x-axis represents the IP address range (which has been mapped to the range 10.0.0.0/24 for privacy concerns). The data depicted in Figure 3 is from a single IPoI over a period of one week. The IPoI was assigned labels for scanning ports 21-23, 51, 80-84, 143. 189 and 443. Based on the pattern observed this particular port scanner was performing a horizontal scan targeting specific ports over the network range. The port scan activity depicted in Figure 3 was observed over a 48 day period. During the period the scanner would scan on average three ports on 16 hosts spread across the /24 IP address space. These scans were so slow and distributed that the organisational firewalls and Intrusion Detection Systems (IDS) did not detect any abnormal activity.

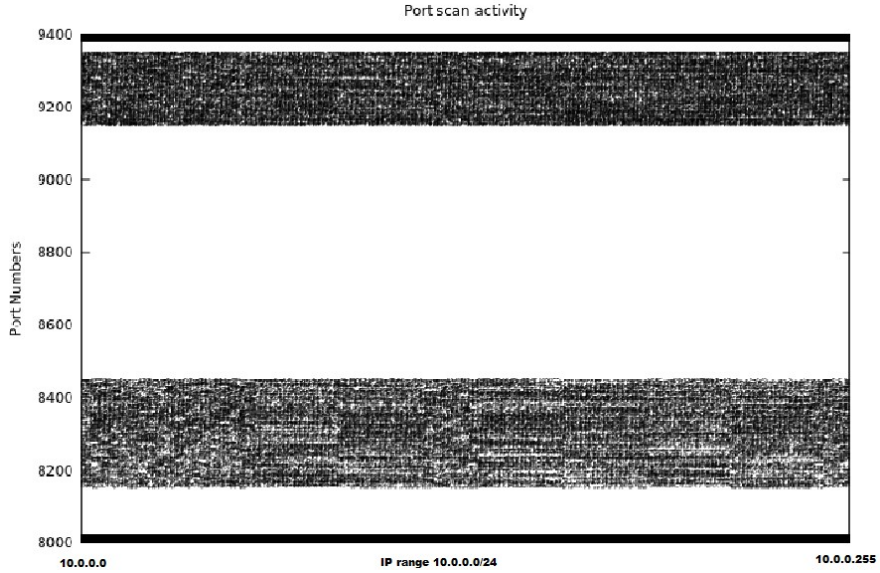

Figure 4: Coordinated block scan activity

In contrast Figure 4 depicts an observation of an aggressive block scan activity. Based on the IPFIX data, the scanner only probed ports which were either between 8151 and 8450 or 9151 and 9350. This specific scanning activity, depicted in Figure 4, formed part of a coordinated scanning effort against one of the tertiary institutions within South Africa. During a 36 hour period, 9 source IPs originating from an internet service provider located in the Netherlands, were performing block scans against the tertiary network. The 9 source

IPs never overlapped scanned port blocks, each had its own dedicated block it was scanning. On average the scanning IPs performed 173 probes per host every 18900 seconds (exactly 5 hours and 15 minutes). The activities observed in Figure 3 and Figure 4 resulted in the source IPs being added to the IPoI database.

The brute-force attack detection algorithms within the SANReN CSIRT is based off of work done by Sperotto et al. (2009) and Hellemons et al. (2012).

Sperotto et al. (2009) suggested that the number of packets observed per network flow (ppf) could be an indicator for the attack phases of a SSH brute-force login attack. By observing a time series of ppf observed during an attack Sperotto et al. (2009) was able to delineate between the pre-attack phase (scanning phase), attack phase and the post-attack phase (residue phase).

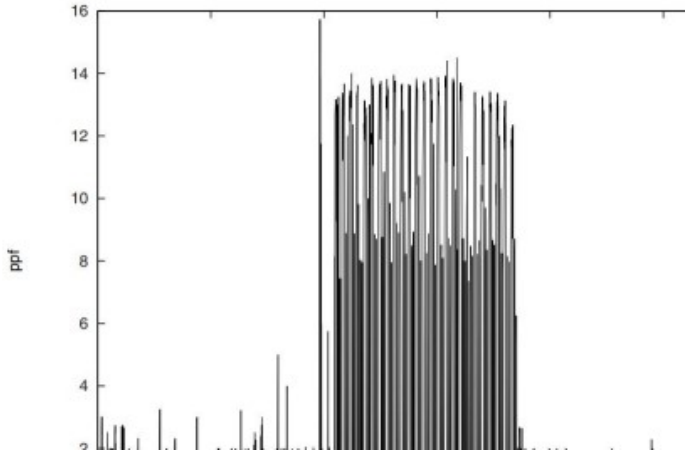

Figure 5: Packets per flow observed during brute force SSH attack Sperotto et al. (2009)

The experimental data used by Sperotto et al. is shown in Figure 5. The scanning phase the majority of ppf measurements are at most 2 and during the brute force attack phase the ppf value fluctuates between 8 and 14 ppf. Sperotto et al. (2009) also noted that between each active brute force attempt a small period of inactivity was observed. By using the data collected during the attack Sperotto et al. (2009) derived the hidden state sequence of the Markov Model. Then by using the hidden sequence the researchers were able to calculate the transition probability between Markov states.

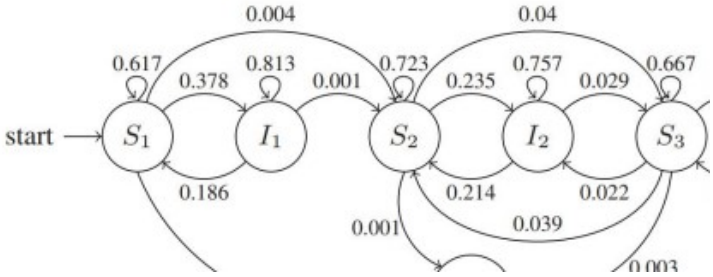

Figure 6: Hidden Markov Model proposed by Sperotto et al. (2009)

The Hidden Markov Model (HMM) to detect SSH attacks, proposed by Sperotto et al. (2009), is shown in Figure 6. S_1, S_2, S_3 represents the three phases of the brute force attack as defined by Hellemons et al. (2012). I_1, I_2, I_3 represents the state of inactivity as observed by Sperotto et al. (2009). It is important to note that based on the proposed model S_1, S_2, S_3 can all proceed to the End state at any point. This implies that an attacker may stop during the attack.

The probabilities shown within the HMM is specific to the attack observed by Sperotto et al. but a similar experimental design was used to define a more generic HMM used by the SANReN CSIRT. The work done by Sperotto et al. (2009) and Hellemons et al. (2012), was also further expanded upon by the SANReN CSIRT to include brute-force attack models for FTP and RDP brute-force attacks.

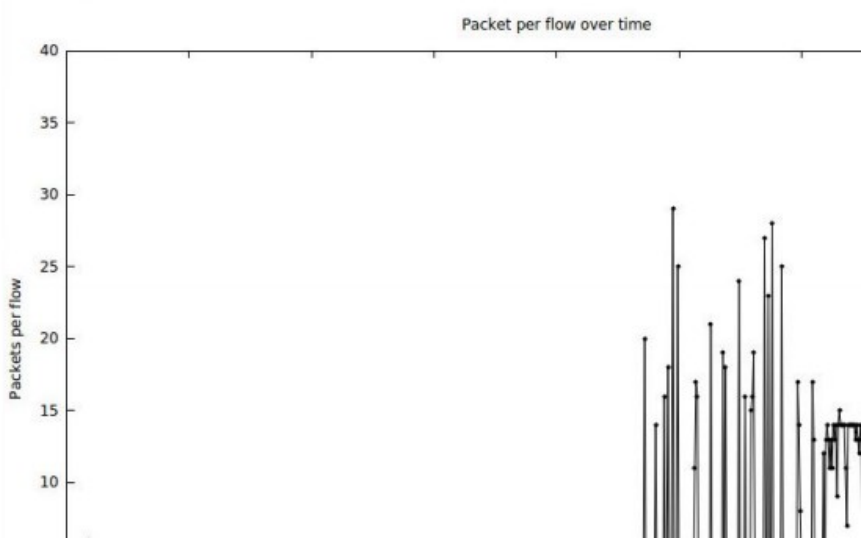

Figure 7: FTP brute-force attack detection within SA NREN

Figure 7 depicts a time series observation of an attempted FTP brute-force login attack against TENET infrastructure components hosted within the SA NREN. The attack pattern is similar to the one observed in Figure 5 by Sperotto et al. During the sensor testing conducted on the NREN network it was determined that fixed using ppf counts as was suggested by Sperotto et al. is not sufficient. Because the SANReN IPFIX collectors are on the TENET backbone the ppf values tended to higher than what was suggested by Sperotto et al. Upon investigation it was determined that security equipment on the SA NREN added additional overhead to the attacker's network flows and as a result the SANReN CSIRT observed higher ppf values. Thus it is recommended that thorough testing and verification of sensor labelling be conducted rather than using the values suggested by Sperotto et al. directly. In order to detect coordinated attacks port scan activity was also considered in our HMM. The HMM thus also needed to be updated to account for attack models which start in S_2 as depicted in Figure 6.

On 28 February 2018, the world's largest Distributed Reflective Denial of Service attack was launched using Memcached servers (Burke et al., 2018). During the attack five SA NREN servers were used to inadvertently attack remote servers. The key findings of the attack using SA NREN is discussed within the SANReN CSIRT's previous publication (Burke et al., 2018), in this paper only the key findings will be highlighted.

	Flows (1.06 Mil)	Packets (34.8 Billion)	Bytes (49.7 Terabyte)	Packets Per Second (276104)	Bytes Per Second (3.1 Gigabyte)
Server 1	38076	1.54E+10	2.21E+13	132005	1.50E+09
Server 2	30470	1.23E+10	1.76E+13	117001	1.30E+09
Server 3	853755	5.50E+09	7.8E+12	1177	1.33E+07
Server 4	13970	1.10E+09	1.6E+12	24616	2.81E+08
Server 5	129528	4.55E+08	6.47E+11	1305	1.49E+07

Figure 8: Summary of SA NREN DRDoS attack, February 2018

During the attack five SA NREN servers were affected by the reflective attack. A summary of the traffic generated between 23 February 2018 and 8 March 2018 is presented in Figure 8. The attack was the first and largest attack detected to date using the IPFIX sensors within the SA NREN. After the attack the following rules were applied to the IPs detected during the attack to label the IPs with the IPoI database:

- If the IP address belongs to a known legitimate user of Memcached service, label the IP as benign. If these benign services form part of an attack, the operators of the services was notified.
- If the number of Memcached requests issued per day increased by more than 50% after 21 February, the IP was labelled as suspicious.
- If the ratio of request traffic to response traffic generated by the IP was greater than 1:100, the IP was labelled as a victim.
- If any of the other network flow sensors, deployed by the SANReN CSIRT, detected the IP as a port scanner, the IP was labelled as a port scanner.
- If any of the port scanner IP addresses were found to only be scanning for port 11211 and using the Memcached protocol, it was assigned an additional label, memcached scanner.

The rule set described above formed the bases for future DDoS attack patterns within the SA NREN.
By combining the observations obtained from the port scan detection algorithms, the brute-force detection filters and the DDoS detection rule sets the SA NREN was able to collect a corpus of label IPoI data. This data was used to train Machine Learning (ML) algorithms to detect potential botnets within the SANReN.
To further enhance the accuracy of the ML models it was supplemented with behavioural models based on patterns observed by common large botnets. For example, as was mentioned in Section 3.2, the Satori botnet Android IoT devices by delivering a malicious payload of ADB (port 5555) in July 2018 (Pour, et al., 2019). Based on the results obtained from the IPoI database it was determined that similar attack behaviour was observed within the SA NREN as early as 5 April 2018.
The SANReN CSIRT built a botnet detection model based on the GoldBrute botnet attack pattern discussed in Section 3.3 and was able to detect similar attack traffic on 11, 15 and 22 September 2019. This despite the GoldBrute C2 servers being taken down in July 2019. The newly observed C2 servers observed to be hosted in Denmark, France, Brazil and Russia.
The SANReN CSIR is currently tracking several variant of Mirai botnets which target IoT devices within the SA NREN. Based on the ports targeted by the botnets and the characteristics of the brute-force attacks used by the bots the ML algorithms are able to cluster the bots into distinct families.

6. Conclusion and future work

In this paper the current efforts of the SANReN CSIRT to detect botnets are discussed. The CSIRT uses a combination of attack sensors to label IPs based on their attack behaviour. This allows the CSIRT to track and label suspicious activity within the SA NREN. At present the majority of the research and detection is inward focused to protect the South African infrastructure. However in future the SANReN CSIR plans to make the IPoI database and label publically accessible for other NRENs to view and contribute.
The botnet components of the SANReN IPFIX attack sensors are still reliant on attack models being built based on known botnet attack models and behavioural analysis. Current research is being conducted within the SANReN CSIRT to develop more generic sensors to detect unknown attack models.

References

Amini, P., Azmi, R. and Araghizadeh, M., 2014. Botnet detection using NetFlow and clustering. Advances in Computer Science: an International Journal, 3(2), pp.139-149.

Antonakakis, M. et al., 2017. Understanding the mirai botnet. In: *26th USENIX Security Symposium USENIX Security 17.*, pp. 1093-1110.

Anubhav, A., 2018. *Masuta: Satori Creators' Second Botnet Weaponizes A New Router Exploit.* [Online]
Available at: https://blog.newskysecurity.com/masuta-satori-creators-second-botnet-weaponizes-a-new-router-exploit-2ddc51cc52a7
[Accessed 5 September 2019].

Bilge, L. et al., 2012. Disclosure: detecting botnet command and control servers through large-scale netflow analysis. *Proceedings of the 28th Annual Computer Security Applications Conference, ACM,* pp. 129-138.

Burke, I.D., Herbert, A. and Mooi, R., 2018, September. Using network flow data to analyse DRDoS attacks, as observed on the SANReN. In Proceedings of the Annual Conference of the South African Institute of Computer Scientists and Information Technologists (SAICSIT), pp. 164-170.

Erquiaga, M. J., Catania, C. & García, S., 2016. Detecting DGA malware traffic through behavioral models. *IEEE Biennial Congress of Argentina (ARGENCON),* pp. 1-6.

Feily, M., Shahrestani, A. & Ramadass, S., 2009. A survey of botnet and botnet detection. *Third International Conference on Emerging Security Information, Systems and Technologies,* pp. 268-273.

GlobalStats, 2019. *Mobile Vendor Market Share South Africa - Oct 2018 - Oct 2019.* [Online]
Available at: https://gs.statcounter.com/vendor-market-share/mobile/south-africa
[Accessed 23 October 2019].

Grégr, M. Portscan detection using netflow data. Proceedings of EEICT10, pp. 229– 233, 2010.

Grill, M., Nikolaev, I., Valeros, V. & Rehak, M., 2015. Detecting DGA malware using NetFlow. *IFIP/IEEE International Symposium on Integrated Network Management,* pp. 1304-1309.

Hellemons, L., Hendriks, L., Hofstede, R., Sperotto, A., Sadre, R., and Pras, A. Sshcure: a flow-based ssh intrusion detection system. In IFIP International Conference on Autonomous Infrastructure, Management and Security, pp. 86–97. Springer, 2012.

Hofstede, R., Bartos, V., Sperotto, A. & Pras, A., 2013. Towards real-time intrusion detection for NetFlow and IPFIX. *Proceedings of the 9th International Conference on Network and Service Management,* pp. 227-234.

Hutchins, M., 2017. *Investigating Command and Control Infrastructure (Emotet).* [Online]
Available at: https://www.malwaretech.com/2017/11/investigating-command-and-control-infrastructure-emotet.html [Accessed 5 October 2019].

Marinho, R., 2019. *GoldBrute Botnet Brute Forcing 1.5 Million RDP Servers.* [Online]
Available at: https://isc.sans.edu/forums/diary/GoldBrute+Botnet+Brute+Forcing+15+Million+RDP+Servers/25002/
[Accessed 30 October 2019].

Mazzariello, C., 2008. IRC traffic analysis for botnet detection. *The Fourth International Conference on Information Assurance and Security,* pp. 318-323.

Pour, M. S. et al., 2019. Data-driven Curation, Learning and Analysis for Inferring Evolving IoT Botnets in the Wild. *Proceedings of the 14th International Conference on Availability, Reliability and Security,* pp. 16-26.

Sweeney, M. & Irwin, B., 2017. A NetFlow Scoring Framework For Incident Detection. *Southern Africa Telecommunication Networks and Applications Conference (SATNAC) ,* pp. 300-305.

Wagner, C., François, J., State, R. & Engel, T., 2011. Machine learning approach for IP-flow record anomaly detection. *International Conference on Research in Networking,* pp. 28-39.

Wijesinghe, U., Tupakula, U. & Varadharajan, V., 2015. An enhanced model for network flow based botnet detection. *Proceedings of the 38th australasian computer science conference ,* Volume 27, pp. 101-110.