SCTE·ISBE CABLE-TEC EXPO 2019 NEW ORLEANS, LA SEPT. 30-OCT. 3

2019 Fall Technical Forum SCTE·ISBE • NCTA • CABLELABS

Predicting the Evolution of Distributed Denial of Service Attacks on Carrier Networks

An Analysis of Shared DDoS Data

A Technical Paper prepared for SCTE•ISBE by

Kyle Haefner
Senior Security Engineer
Cable Television Laboratories Inc.
858 Coal Creek Circle, Louisville, CO 80027
Phone:303-661-3320
k.haefner@cablelabs.com

Table of Contents

Title **Page Number**

Table of Contents 2
Introduction 4
Background 4
1. DDoS Information Sharing Project Overview 4
2. DDoS Introduction and History 5
2.1. Cost of DDOS 5
3. Taxonomy of DDoS Attacks 5
3.1. Overview of DDoS Detection 6
Methods and Results 6
4. Global Attack Statistics 7
5. Attack Prediction 9
5.1. Long Term Short Term Memory (LSTM) Attack Prediction 9
6. Shodan Data 10
6.1. Random Forest Predictions on Shodan Data. 12
Conclusion 15
Appendix 16
Abbreviations 16
Bibliography & References 17

List of Figures

Title **Page Number**

Figure 1: DDoS Information Sharing Overview 4
Figure 2: Top Attacks by Type 7
Figure 3: Top Attacks Countries 8
Figure 4: Top Attacks by Bandwidth 8
Figure 5: Top Attacks by Packets 8
Figure 6: Top Attacks by Duration 9
Figure 7: LSTM Attack Prediction 10
Figure 8: Top Attackers by Service 11
Figure 9: Top Attackers by Protocol 11
Figure 10: DNS Amplification Top Features 12
Figure 11: NTP Amplification Top Features 13
Figure 12: CLDAP Amplification Top Features 13
Figure 13: IP Fragmentation Top Features 14
Figure 14: Total Traffic Top Features 14
Figure 15: UDP Top Features 14

List of Tables

Title **Page Number**

Table 1: Top attackers by operating system 11
Table 2: Classifier results on attack data 12
Table 3: Attack Labels Used 16

Introduction

Distributed Denial of Service (DDoS) attacks are among the preeminent threats facing the Internet today. Predicting where the next DDoS attack will emanate at an endpoint/subscriber level is a long-sought goal of the cyber-security community.

This work evaluates attack data from five contributing members of the DDoS Information Sharing (DIS) project with the intent to provide an ISP/MSO the tools to predict ***at subscriber/endpoint granularity*** if they will start participating in a DDoS attack. The DIS data is combined with data from the Internet search engine, Shodan, to build a detailed dataset of recent/active attackers. Statistical and machine learning analysis of this composite dataset demonstrates that by evaluating network endpoints with certain features, it can be predicted that these endpoints will participate in a specific type of DDoS attack with accuracies between 91-98%.

Finally, each feature of the attacking network endpoint that was used in the machine learning model is ranked by its predictive significance, lending insight into how ISP/MSOs might *preemptively* detect and mitigate an endpoint even before it starts participating in a DDoS attack.

Background

1. DDoS Information Sharing Project Overview

The DIS project began as a pilot in early 2017 through the Messaging, Malware, Mobile Anti-Abuse Working Group (M3AAWG). At its core, its purpose is to allow large ISPs to share their attack data with the goal to help them remediate compromised and vulnerable systems running within their own networks.

This is accomplished by using a trusted third party to aggregate attack data from participants and provide API access to this aggregated data in a way that the ISP can see the attacks that are emanating from within their own AS (Autonomous System) as shown in Figure 1.

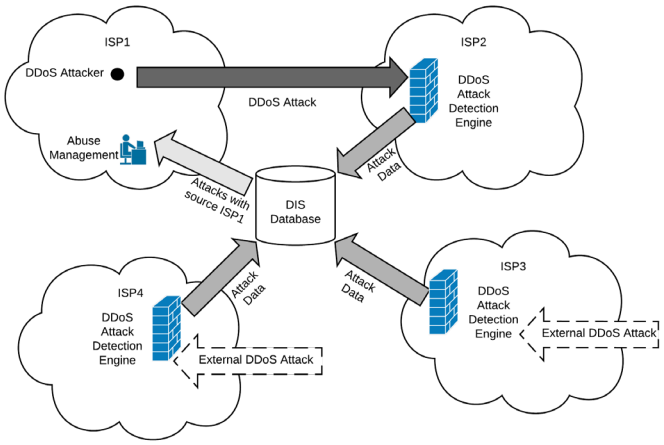

Figure 1: DDoS Information Sharing Overview

2. DDoS Introduction and History

Distributed Denial of Service attacks are attacks that emanate from many different and distributed sources and usually target a single entity on the Internet. The goal of the attacker is to overwhelm a service provider and deny legitimate users' access to the service.

The first notable DDoS attack that was widespread occurred in February 2000 by a 15-year-old going by the handle MafiaBoy (Michael Calce) (HERSHER, 2015). MafiaBoy enlisted several University servers to send many simultaneous requests and brought down some of the biggest names in e-commerce at the time including Amazon, CNN, Dell, eBay and Yahoo. The attack itself was small in scale and simple by today's standards, using only a handful of powerful computers to submit many legitimate requests simultaneously. The webservers of the time were overwhelmed and unresponsive for a span of several hours and in some cases several days.

In 2007 one of the first documented cyber-warfare DDoS attacks crippled the government of Estonia (Goth, 2007). The attack was relatively small in terms of bandwidth and targeted only a handful of websites however the collateral damage overwhelmed Estonia's network infrastructure effectively taking the entire nation offline. The politically motivated attack is largely blamed on Russian actors and resulted in the drafting of new international laws, notably the Tallinn manual on the international law applicable to cyber warfare (Schmitt, 2013).

The Mirai botnet attack of 2016 was not only one of the largest DDoS attacks in terms of bandwidth at 1.1 Tbps it, was also one of the most disruptive (Kolias, 2017). Composed of a botnet of over 600k compromised Internet-of-Things (IoT) devices, this attack targeted high-profile services such as the Dyn DNS (Domain Name Service) provider blocking access to many popular websites, such as Twitter, Netflix, Reddit and GitHub for many hours.

As of this writing the largest ever published DDoS attack in terms of bandwidth was reported by Imperva, where in April of 2019 they reported a SYN DDoS attack of 500 million packet-per-second attack resulting in a phenomenal 3.4 Tbps (Crane, 2019)!

2.1. Cost of DDOS

Denial of service attacks are so effective because they are extremely cheap for the attacker and extremely expensive for the victim. The rise of DDoS as a Service (DDoSaaS) on the dark web has commercialized and commoditized these types of attacks, drastically lowering the total cost and barrier to entry required to launch them. An analysis by Kaspersky Labs examined several DDoSaaS providers on the dark web and found that a DDoS attack lasting 300 seconds with a bandwidth of 125 Gbps will cost as low as $6 (US). Others advertise an hourly rate of $20 per hour for attacks in the hundreds of Gbps, and offer various plans and a simple pricing structure based on type and scope of attack (Makrushin, 2017).

For victims the cost of a DDoS is much higher, Incapsula surveyed 270 North American organizations and estimated that a targeted DDoS attack costs a victim an average of $40,000 per hour (Mathews, 2014). B2B International research firm estimated that a DDoS attack costs enterprises an average of $2 million per incidence (Kobialka, 2018).

3. Taxonomy of DDoS Attacks

DDoS attacks can generally be divided into three broad categories, volumetric attacks, protocol specific, and application specific. The DIS data have examples of attacks from all three categories.

Volumetric Attacks

Volumetric attacks are designed to saturate the target network by flooding it with traffic. Some examples of this type of attack include UDP floods, and ICMP floods. Volumetric attacks often have a broad affect even if they have a narrow target. For example, a flood of spoofed UDP traffic aimed at specific server could saturate many of the network paths leading to that server causing access to other websites and services to become unreachable.

Protocol Attacks

Protocol attacks take advantage of specific weakness in protocols and focus on depleting resources. Protocol attacks work against specific servers, or intermediary network equipment such as firewalls, load balancers and SDN controllers.

Application Attacks

This type of attack does not produce high levels of network traffic, instead the attack is targeted at specific server applications with the goal of making the service exhaust CPU or memory resources.

3.1. Overview of DDoS Detection

Detecting DDoS attacks is a challenging problem due to the heterogenous nature of how these attacks are carried out and requires an equally heterogenous set of solutions. There are three main categories that are used for DDoS detection, statistical methods, knowledge-based methods and machine learning based methods. Often combinations of these three are used simultaneously.

Statistical Methods

Statistical methods generally attempt to model the normal traffic and then test any new traffic or flows to determine if it belongs to the normal set or is an anomaly.

Knowledge Based Methods

Knowledge based approaches to detecting attacks use predefined rules and patterns to determine if the traffic is an attack or not. Some examples of knowledge-based approaches can be as simple as threshold-based systems to more complex state-transition and signature analysis.

Machine Learning Methods

Machine learning methods of detecting DDoS attacks apply machine learning enormous ability to absorb vast amounts of data and learn classifications across that data.

Methods and Results

This section is divided into two parts. Part one analyzes global statistics of attacks across all participants, giving a broad picture of the scope, scale and structure of attacks. A prediction algorithm that uses sequential data was run on the global data to predict the next attack in the sequence.

Part two uses recent attack events from the DIS project to query Shodan, a search engine for connected devices. The Shodan query returns fine grained details such as port, service and operating systems of

active/recent attackers. These details are then composed into features and fed into machine learning predictive models.

4. Global Attack Statistics

Global statistics are generated from attack events submitted across all DIS participants. Attack events are purged by the system if they are older than 30 days; however, reports are generated on a weekly basis and sent to participants. These weekly reports store aggregated meta data such as top attack types, countries etc. The following graphs show data from the first six months of 2019.

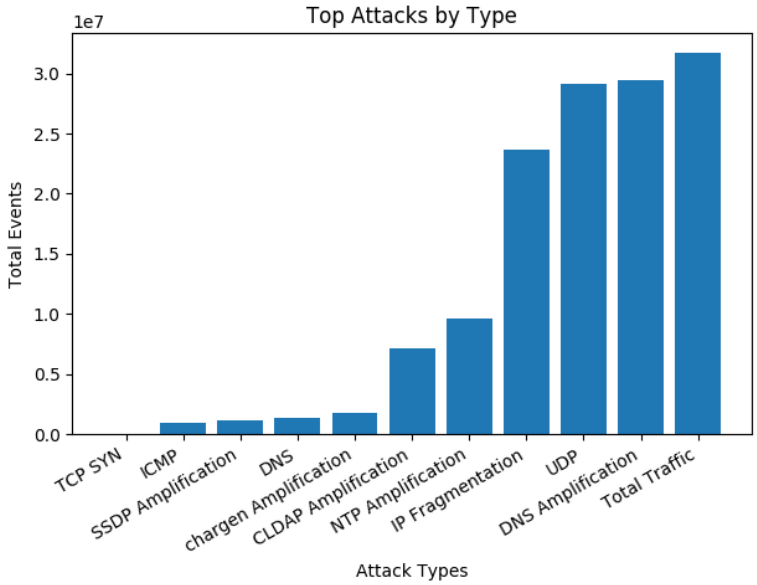

Figure 2: Top Attacks by Type

Figure 2 shows the top attack types seen by the DIS system over the first six months of 2019.

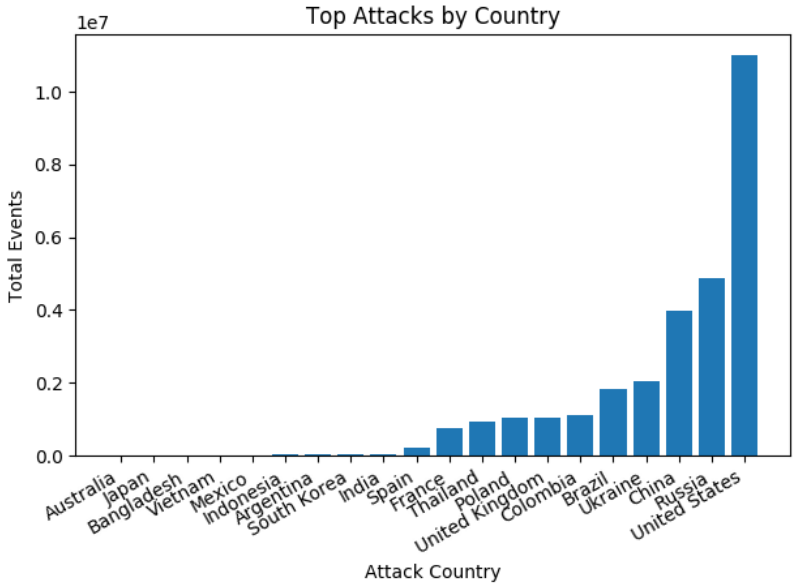

Figure 3: Top Attacks Countries

Figure 3 show the origin of attacks by county for the first six months of 2019. These countries and their rank are consistent with other sources and analysis. (Link) (Akamai, 2019)

The following figures represent a sampling of data gathered between June 1, 2019, and June 30, 2019.

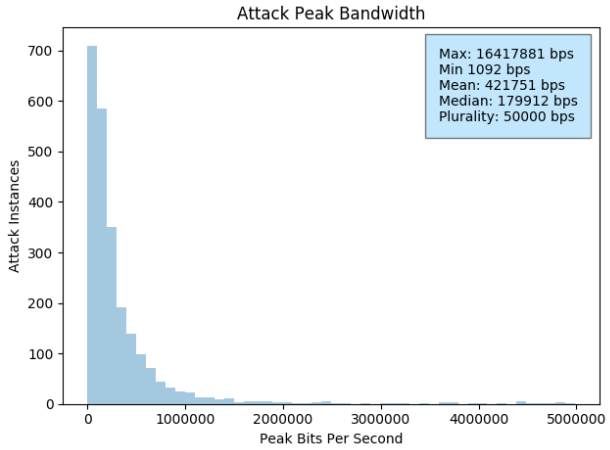

Figure 4: Top Attacks by Bandwidth

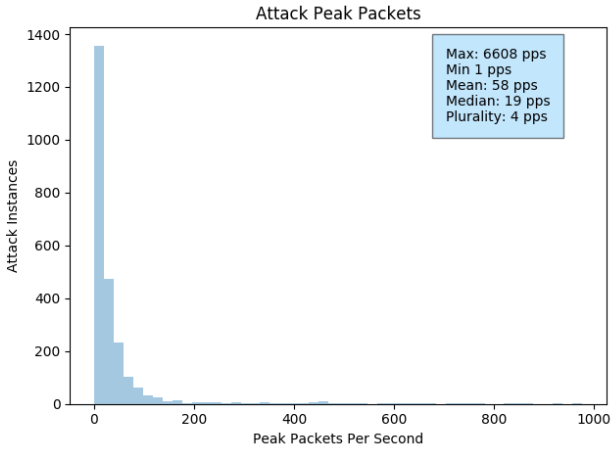

Figure 5: Top Attacks by Packets

Figure 4 above indicates that most of the attacks are not volumetric in nature with the median attack only 179Kpbs and the plurality (mode) of the attacks having a bandwidth of only 50Kbps.

Figure 5 furthers this observation where the median packets per second is only 19 pps and the plurality (mode) is only ***4 pps***.

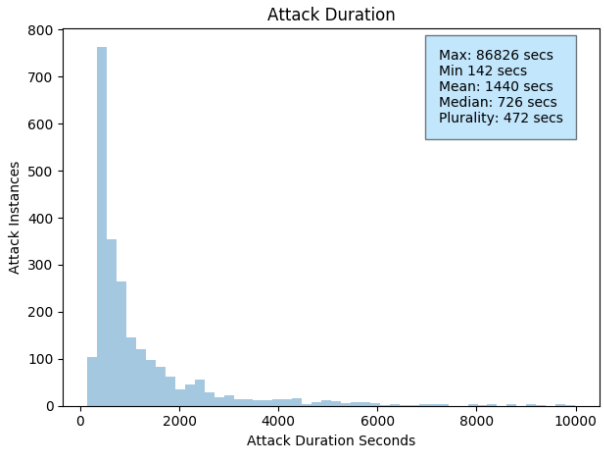

Figure 6: Top Attacks by Duration

Figure 6 shows that the average attack is just over 20 minutes and the most frequent attack lasts only a little more than 7 minutes. Figure 6 also shows that the longest measured attack was almost 60 *days* long (active and ongoing attacks are not purged every thirty days).

5. Attack Prediction

Attack data from participants comes with a label (Table 3) generated by the Netscout Arbor system running on the participant's network. This global data only contains a few usable features making it largely unsuitable for predictive analysis. However, some predictive analysis can be run using only the labels and the associated IP address.

5.1. Long Term Short Term Memory (LSTM) Attack Prediction

Recurrent Neural Networks (RNNs) are machine learning neural networks that are useful on time-series data. Long-Term Short-Term Memory (LSTM), are subset of RNNs that can learn long term dependencies and are particularly well suited to learning from large sets of sequential data and are used extensively on word prediction algorithms.

The hypothesis in this analysis is that a host that is compromised by one malware is likely to become compromised by another malware. Each malware forms a botnet that propagates unique types of attacks. These attacks form a timeline of attacks such that a host that propagates attack type A, later propagates attack type B. The predictive model analysis then follows that if a $Host_X$, that propagates a series of attacks,

$Host_X$: $Attack_A \rightarrow Attack_B \rightarrow Attack_C \ldots$

how likely is that host to also exhibit $Attack_D$?

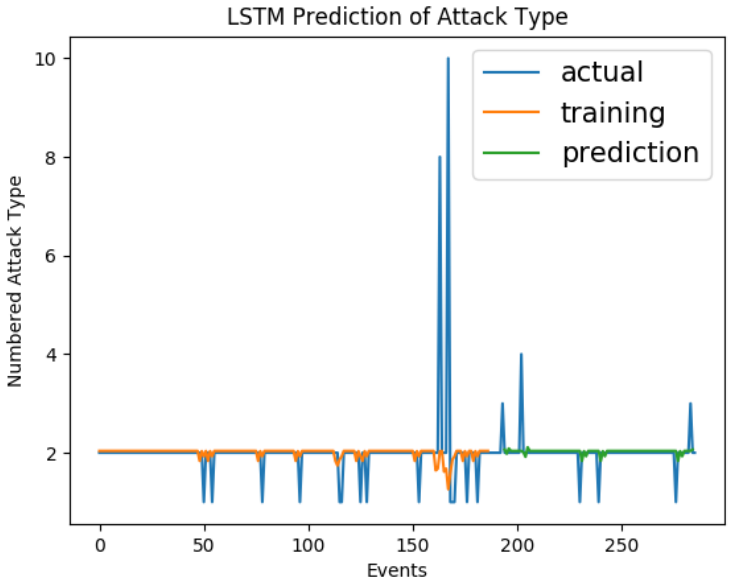

Figure 7: LSTM Attack Prediction

As can be seen in Figure 7, the blue line is the series of attack data from a single IP address that propagated several types of attacks (in order for the model to analyze the data, attacks names were each converted a number see Table 3: Attack Labels). The orange line represents the model being trained and the green line is the model's prediction.

Based on the graph, the LSTM prediction model did *not* accurately predict the attack type. The primary reason for an LSTM model to fail is that the data is too random in nature to be predicted. To confirm this hypothesis, the augmented Dickey-Fuller (ADF) test was run across the data (Dickey, 1979).

The ADF test shows how strongly series-based data can be defined by a trend. If the ADF test results in a positive score, then the data has a series-dependent structure to it. If the score is negative, then the data is too random to predict. The ADF score for the DIS labeled data run across several hosts that had multiple attack types resulted in a score of -5.9611. This score implies that the data is too random to perform time-series predictions.

Why is the data random? It is believed that the DHCP assignment of IP addresses to attacking hosts introduces noise into the data that prevents any predictive analysis to be run. The LSTM model could be a viable method of prediction if host data could be statically tied back to a consistent IP address. For MSO/ISPs, this could be accomplished with internal hosts using current DHCP assignment logs.

6. Shodan Data

Shodan is a search engine that continuously scans the Internet for open and accessible ports. When Shodan successfully connects to a port it grabs the response from the host and stores and indexes the returned data, called the banner, along with the IP address in a database. Shodan offers an application program interface (API) to query this database.

For this research the most recent 100K events representing recent/active DDoS attacks from the DIS service were used. For each event the IP was extracted and then used to query the Shodan database. Queries that returned a recent record from Shodan are considered to be recent/active attackers.

Of the 100K events in DIS, Shodan returned 2,388 recent records. By parsing the banner of each returned record, and transforming the categorical features (operating system, protocol, etc.) into numbers using a technique called one hot encoding, the data was expanded to 231,226 examples (rows) and 225 features (columns). From this new composite dataset some basic aggregation analysis was done across the dataset. Table 1 below shows a breakdown of the attackers based on operating system.

Table 1: Top attackers by operating system

Operating System (OS)	Attacker Count
Windows Server (various versions)	201
Windows Desktop (various versions)	52
Linux 3.x	48
Windows Embedded	6
Linux 2.6.x	4
Darwin (MacOS)	3

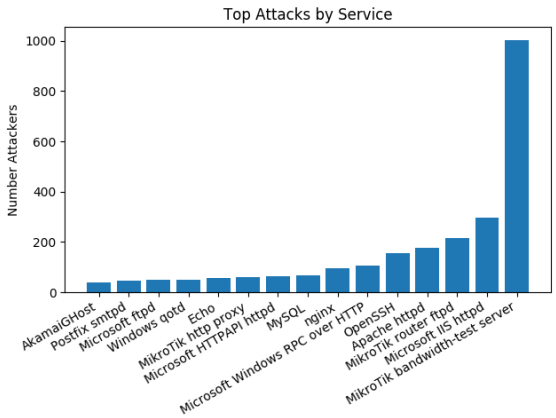

Figure 8: Top Attackers by Service

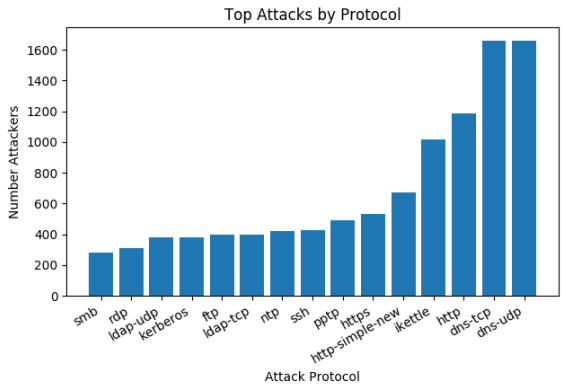

Figure 9: Top Attackers by Protocol

Figure 8 shows the breakdown of attackers based on service running on the host. It is notable that the MikroTik bandwidth test service is detected at a rate nearly five times more than the next nearest service, Microsoft's IIS http (web) server.

Figure 9 shows the breakdown of attackers by protocol. Each of the protocols was derived from the Shodan module that was used to detect it. Domain Name Service , a frequent source of DDoS attacks due to the high amplification potential of the protocol, takes the top two slots.

The appearance of the ikettle protocol in the top four was a notable result. The iKettle protocol is a binary protocol that runs over UDP or TCP and port 2081. This protocol is used to control a smart WIFI-connected kettle of the same name used to heat water for tea or coffee. iKettles come with a default password of "000000". Once the iKettle is joined to the user's network, it allows connections to port 23 using Telnet. An attacker can connect to the iKettle and request it to list its settings, one of which is the WIFI password that is in plain text (hughes, 2015).

6.1. Random Forest Predictions on Shodan Data.

The composite dataset has the label imported from the DIS data and many additional features from the Shodan data, so supervised learning classifiers can now be deployed to predict the type of attack using the weighted Random Forest classifier (Liaw, 2002). The weighted Random Forest classifier belongs to a set of classifiers that use an ensemble of decision trees to build a model of the data.

The weighted Random Forest classifier works by building several predictive models and votes on the best one. This classifier is particularly useful when applied to unbalanced data, i.e., where the labels are not evenly split. The dataset derived from top six DIS attack types as labels combined with the Shodan host data had majority sets made up of attack types DNS amplification and NTP amplification with the other four belonging to minority sets. The weighted Random Forest model assigns a higher weight and misclassification cost to the minority classes this in turn reduces the bias toward the majority classes.

The composite attack dataset was split into two parts where 80% was used for training and 20% was reserved for testing. Table 2 shows the accuracy of the classifier of predicting the correct label on the test data.

Table 2: Classifier results on attack data

Attack Label	Random Forest Classifier Accuracy
DNS Amplification	98.036%
NTP Amplification	98.468%
CLDAP Amplification	95.055%
IP Fragmentation	91.204%
Total Traffic	96.265%
UDP	98.234%

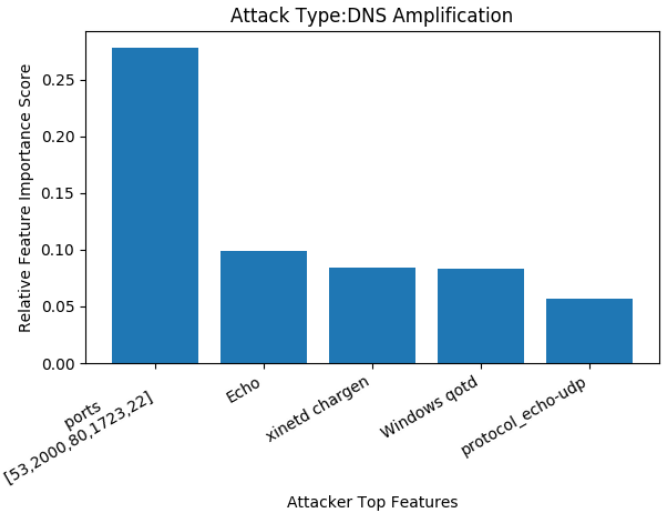

Figure 10: DNS Amplification Top Features

Figure 10 shows the top features used in predicting a DNS amplification attack, as would be expected port 53 is seen. Echo refers to the echo protocol (RFC 862) that is associated with the init.d services on Linux. Based on Figure 10, this classifier utilized the ports feature predominantly in predicting the attack classification of DNS amplification. QOTD is a quote of the day service.

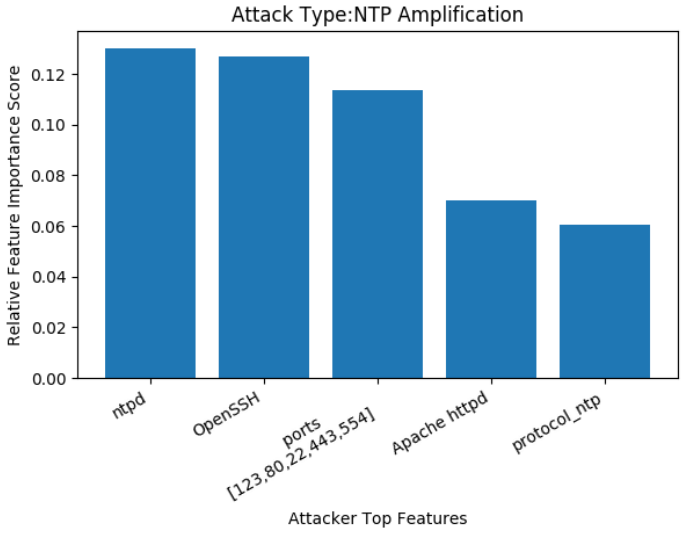

Figure 11: NTP Amplification Top Features

Figure 11 shows the features for an NTP amplification attack. Features for the ntpd service, ntp protocol and port 123 make logical sense here. The figure shows a close ranking of feature importance, especially with the ntpd service, OpenSSH and the ports features. This implies these features are of relative equal importance to the attack classification of NTP amplification.

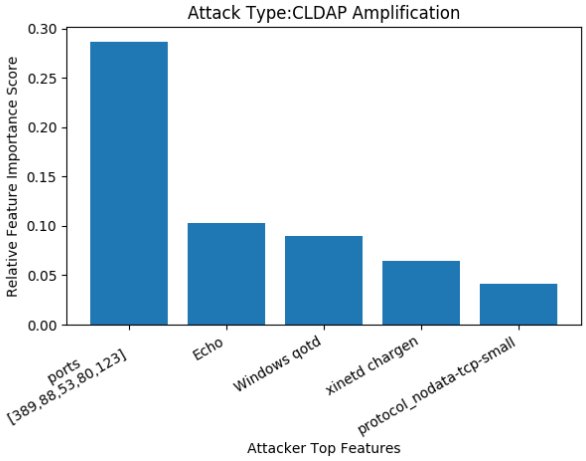

Figure 12: CLDAP Amplification Top Features

Figure 12 shows the features used in the CLDAP classification. CLDAP is associated with Microsoft Active Directory services and often serves as an application layer 'ping' for Active Directory. Here the open ports were the predominant feature used by the classifier.

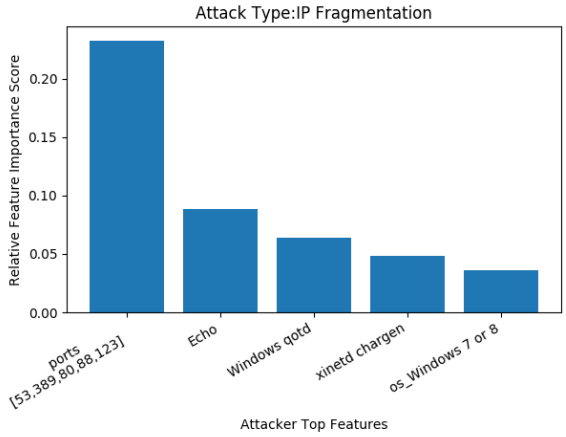

Figure 13: IP Fragmentation Top Features

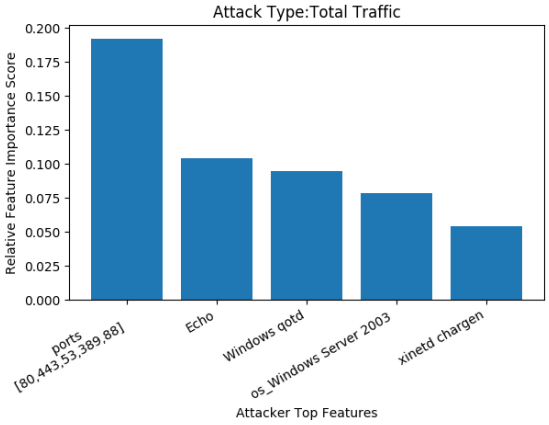

Figure 14: Total Traffic Top Features

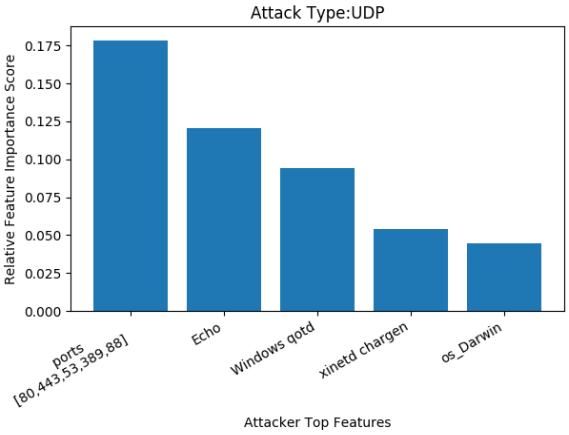

Figure 15: UDP Top Features

Figure 13, Figure 14, and Figure 15 are grouped together as they have similar predominant features. Despite having similar features, the classifier was able to predict each attack classification with a greater than 90% accuracy rate.

Conclusion

In this research, DDoS Information Sharing (DIS) data was explored showing the top attack types and top attack origins as seen by the participating members. Using the sequence of attacks from each IP address, the feasibility of using LSTM RNNs for predicting the next attack type in the sequence for a given IP address was examined. The current data is too random to pursue this method by itself, however, it is believed with the addition of DHCP lease logs, this remains a viable prediction model.

Next, the latest events seen by the DIS service was used to query the Internet search engine, Shodan. From these queries, a detailed dataset was built for recent/active attackers that showed the top operating systems, services and protocols running on each attacker.

This new dataset was analyzed using the Random Forest ensemble classifier to predict the attack type of an endpoint based on open ports and the information these hosts present when connecting to them. The model was able to correctly predict the DDoS attack type with accuracies above 90% for each type of attack. Lastly, a breakdown of each attack type and the features of the attacker that are most important to the predictive model was shown.

The research presented in this paper could be directly applied by an ISP/MSO to predict which subscribers have a node that have been or can be compromised on their network. This in turn could be used in remediation efforts and upstream DDoS prevention services, preemptively, before the compromised node has started participating in a DDoS attack.

Appendix

Table 3: Attack Labels Used

#	Attack Name	#	Attack Name
1	UDP	12	NTP Amplification
2	Total Traffic	13	SSDP Amplification
3	DNS Amplification	14	chargen Amplification
4	TCP RST	15	SNMP Amplification
5	IP Fragmentation	16	MS SQL RS Amplification
6	DNS	17	rpcbind Amplification
7	TCP SYN	18	memcached Amplification
8	TCP SYN/ACK Amplification	19	RIPv1 Amplification
9	ICMP	20	mDNS Amplification
10	CLDAP Amplification	21	NetBIOS Amplification
11	TCP NULL		

Abbreviations

API	application program interface
AS	Autonomous System
bps	Bits per second
CLDAP	Connection-less Lightweight Directory Access Protocol
DDoS	Distributed Denial of Service
DDoSaaS	DDoS as a Service
DHCP	Dynamic Host Configuration Protocol
DIS	DDoS Information Sharing
DNS	Domain Name System
Gbps	Gigabits per second
ICMP	Internet Control Message Protocol
IoT	Internet of Things
Kbps	Kilobits per second
LSTM	Long-Term Short-Term Memory
M3AAWG	Messaging, Malware, Mobile Anti-Abuse Working Group
NTP	Network Time Protocol
RNN	Recurrent Neural Network
SDN	Software-Defined Networking
Tbps	Terabits per second
TCP	Transmission Control Protocol
UDP	User Datagram Protocol

Bibliography & References

Akamai. (2019, July 06). *Web Attack Visualization*. (Akamai) Retrieved July 14, 2019, from https://www.akamai.com/us/en/resources/our-thinking/state-of-the-internet-report/web-attack-visualization.jsp

Crane, C. (2019, May 29). *The Largest DDoS Attacks in history.* Retrieved July 1, 2019, from Casey Crane Read more at: https://www.thesslstore.com/blog/largest-ddos-attack-in-history/

Dickey, D. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American statistical association, 74*(266a), 427-341.

Goth, G. (2007). The Politics of DDoS attacks. *IEE Distributed Systems Online, 8*(8), 3-3.

HERSHER, R. (2015, February 7). *Meet Mafiaboy, The 'Bratty Kid' Who Took Down The Internet.* Retrieved July 1, 2019, from https://www.npr.org/sections/alltechconsidered/2015/02/07/384567322/meet-mafiaboy-the-bratty-kid-who-took-down-the-internet

hughes, M. (2015, October 23). *Why the iKettle HAck Should Worry You (Even if You Don't Own One).* Retrieved July 14, 2019, from https://www.makeuseof.com/tag/ikettle-hack-worry-even-dont-one/

Kobialka, D. (2018, February 25). *Kaspersky Lab Study: Average Cost of Enterprise DDoS Attack Totals $2M.* Retrieved July 1, 2019, from https://www.msspalert.com/cybersecurity-research/kaspersky-lab-study-average-cost-of-enterprise-ddos-attack-totals-2m/

Kolias, C. (2017). DDoS in the IoT: Mirai and other botnets. *Computer, 50*(7), 80-84.

Liaw, A. a. (2002). Classification and regression by randomForest. *R News, 2*(3), 18-22.

Link, C. (n.d.). *CenturyLink 2018 Threat Report* . Century Link.

Makrushin, D. (2017, March 23). *The cost of launching a DDoS attack.* Retrieved July 1, 2019, from https://securelist.com/the-cost-of-launching-a-ddos-attack/77784/

Mathews, T. (2014). *Incapsula Survey: What DDoS Attacks Really Cost Businesses.* Retrieved July 1, 2019, from https://lp.incapsula.com/rs/incapsulainc/images/eBook%20-%20DDoS%20Impact%20Survey.pdf

Schmitt, M. N. (2013). *Tallinn manual on the international law applicable to cyber warfare.* Cambridge University Press.

