Impact-based optimisation of BGP Flowspec rules for DDoS attack mitigation

Diederik Bakker
University of Twente
P.O. Box 217, 7500 AE Enschede
The Netherlands
d.r.bakker@student.utwente.nl

ABSTRACT

Distributed Denial of Service (DDoS) attacks aim to take Internet services offline and the frequency and scale of these attacks is increasing. BGP Flowspec defines a protocol to rapidly deploy rules consisting of filters and actions on Internet traffic. Related research shows its potential for DDoS attack mitigation. The protocol allows for taking action on large volumes of traffic, but impact to end-users is imminent because of its low granularity in rule specification. In this paper, we provide a method for quantifying end-user impact of BGP Flowspec rules, including a practical solution to deploy rules into the network. The goal of this research is reducing end-user impact while mitigating an ongoing DDoS attack using BGP Flowspec.

Keywords

DDoS, BGP Flowspec, flow specification, impact quantification, rule generation, rule optimisation

1. INTRODUCTION

Modern society is increasingly dependent on the Internet, and this dependency is only growing. We depend on the internet for communication, keeping up with the news, banking and reading this paper. One of the things standing in the way of our desire for a dependable Internet, is the increasing frequency and scale of Distributed Denial of Service (DDoS) attacks. A DDoS attack has the goal of taking a specific user or service offline or at least reduce its availability for its intended users. Detection and mitigation of these attacks is a very hard challenge because of their distributed nature.

Compared to previous years, the first quarter of 2018 showed an increase in the amount of DDoS attacks [1]. In February 2018, the biggest DDoS attack was recorded at 1.3Tb/s [2]. Santanna et al. showed that DDoS attacks are easy to execute with DDoS-as-a-Service Providers, even for people without technical knowledge [3]. The ease and scale of DDoS attacks show an ongoing problem which calls for an effective solution.

One of the possible solutions for this problem is recognising these attacks and storing a pattern that uniquely identifies the attack, also referred to as a DDoS attack fingerprint.

30th Twente Student Conference on IT Febr. 1st, 2019, Enschede, The Netherlands.

Service providers and enterprises can fingerprint and share traffic that passes through their network, and in turn block the traffic that matches known attack fingerprints from others. The DDoSDB project is an example which uses this approach, and it provides both a standard for defining attack fingerprints and a distribution medium [4].

There are various stages where a DDoS attack can be mitigated, one of them being the Internet Exchange Point (IXP). At this stage, traffic is routed between different autonomous systems of Internet Service Providers (ISP) using the Border Gateway Protocol (BGP). When DDoS attacks are mitigated at this level, it is often performed by an ISP network administrator who manually defines and deploys BGP routes. This method, also called remotely triggered blackholing (RTBH), is time-consuming, error-prone does not allow for very granular traffic filtration.

A recent development is the creation and adoption of BGP Flowspec [5], which is an extension to the BGP protocol. BGP Flowspec offers the ability to create rules on traffic flows and apply corresponding actions. The protocol can play an important role in the mitigation of a DDoS attack, given that the correct rules are deployed in the network. Although it is a large improvement over RTBH, BGP Flowspec still does not allow for very high granularity in its rule specification (table 1), which creates an interesting trade-off. On the one hand, it is possible to discard large amounts of traffic at this level, and prevent nodes deeper into the network from congesting and the attack succeeding. On the other hand, this might come at a cost, because it is very easy to discard normal traffic. This would result in non-malicious traffic being discarded, potentially causing normal users to experience impact.

For this research, we will first look at the possibilities BGP Flowspec offers to specify rules on traffic, as well as its limitations. We then present an algorithm allowing automatic generation of potential BGP Flowspec rules based on known DDoS attack fingerprints. We will then analyse the generated rules with the goal of measuring their impact in the network, both on malicious and non-malicious traffic. This mechanism allows for the selection of the most effective rules. These rules can be presented to an ISP network administrator, who then has the ability to pick and deploy the rules in the network. These contributions will allow for deployment of BGP Flowspec to effectively mitigate DDoS attacks while reducing end-user impact to a minimum.

This paper starts with an overview of related work in section 2. This is followed by section 3, which gives an introduction of all data sources that are used in this research. Section 4 presents a method of generating BGP Flowspec rules from DDoS attack fingerprints. Following that, section 5 will present an algorithm that quantifies

end-user impact and effectiveness of BGP Flowspec rules. Finally, section 6 demonstrates the complete solution using a simulation environment in which the results have been rendered. The paper is finalised with a conclusion and a discussion of future work in sections 7 and 8.

2. RELATED WORK

Hinze et al. [6] look into the potential of BGP Flowspec in relation to more traditional methods of discarding traffic at inter-domain level like remotely triggered black-holing (RTBH). Using RTBH, all traffic to the attack destination would be dropped, causing high impact to normal users as the service would effectively be offline. The benefit of deploying BGP Flowspec, even with little additional information, is shown. However, generation of rules is still a manual job, and impact is only taken into account after deploying the rules in the network. They conclude that IXP-level DDoS mitigation using BGP Flowspec has high potential.

Van Gijtenbeek and Dijkhuizen [7] propose an IXP-level DDoS mitigation method based on RTBH. Their method incorporates network administrator intervention, where manual control is taken to initiate detection and mitigation. The administrator manually chooses destination prefixes to apply RTBH rules. Although we propose a solution where less manual work is required because rules are generated and presented to the administrator, we do see the value in manual intervention before rules are deployed into the network.

Loibl and Bacher [8] created an experiment with carrier aggregation router hardware from different vendors, simulating an inter-AS (IXP) environment. They show that router vendor implementations of the BGP Flowspec standard [5] contain multiple bugs and are missing features. They render it unsafe to use in inter-AS environments in its current state because of BGP sessions terminating upon propagation of certain rules, causing complete network failure which can be triggered remotely. The bugs that were found have been reported with the router vendors. Most of the reported problems are of temporary nature because the underlying issues reside in the implementation of the standard, but more structural limitations exist as well. These limitations should be taken into account when generating BGP Flowspec rules and are discussed in more detail in section 4.2.2.

Steinberger et al. [9] have conducted a survey under Internet Service Providers and other network operators in 2015. Part of this survey was a question about their technical ability to use BGP Flowspec, which 48% of the respondents had. Network hardware vendors have improved support ever since, and therefore these numbers might currently be higher. This shows that BGP Flowspec-based DDoS mitigation solutions should be ready for deployment in a real-world environment.

3. DATA SOURCES

Multiple data sources will be used in this research, allowing the creation and verification of the rule generation and impact quantification algorithms that will be presented in this paper. This section will present a short overview of each dataset and its characteristics.

3.1 DDoS attack data

The dataset used for DDoS attack data is extracted from DDoSDB [4], a platform that helps DDoS attack victims and the academic community to get access to information about DDoS attacks. DDoSDB collects its data from collaborators, often DDoS attack victims, that collected data from DDoS attacks they experienced. After anonymising the data to protect the identity of the victim, it is published on the DDoSDB website in the form of a traffic sample and a fingerprint, which summarises all relevant characteristics of the DDoS attack in a standardised format.

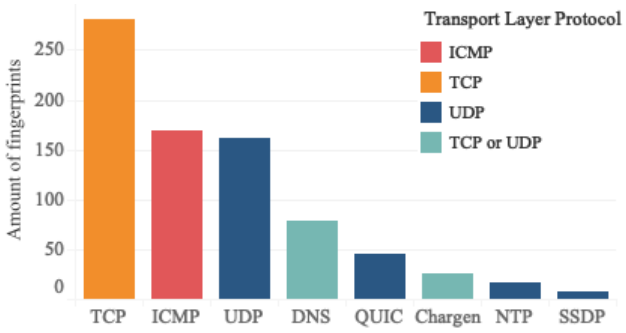

Figure 1. Distribution of protocols in DDoSDB attack fingerprints, coloured by their corresponding transport layer protocol

As of January 2019, the DDoSDB dataset consists of 862 DDoS attack fingerprints and their corresponding trace files. The fingerprints contain 71 protocols from different layers.

We are only interested in protocols that work on the transport layer, because these can be used as a matching criterion in BGP Flowspec rules. Manual mapping of DDoS attack fingerprints onto their corresponding transport layer protocol is very time consuming, and most protocols occur less than 5 times in the dataset. Therefore, the decision has been made to filter out all protocols that occur less than five times in the dataset. This leaves 785 fingerprints, 91% of the complete dataset.

Figure 1 shows the resulting dataset, categorised by the protocol name that is included with the DDoSDB attack fingerprint. These protocols have been coloured by their underlying transport layer protocol. Some higher-level protocols, like Chargen and DNS, can be used with multiple protocols on the transport layer.

3.2 Traffic data

In order to quantify impact of generated BGP Flowspec rules, we need a representative capture of real-world network traffic. TCPReplay provides a capture file that contains "real network traffic on a busy private network's access point to the Internet" [10]. The capture file, called *bigFlows.pcap*, contains 40686 traffic flows using 132 different network protocols, for a total of 791615 packets.

4. BGP FLOWSPEC RULE GENERATION

The DDoSDB fingerprints allow generation of multiple BGP Flowspec rules that have the potential of (partially) mitigating the DDoS attack. This section provides better understanding of the process of generating candidate BGP Flowspec rules.

4.1 BGP Flowspec rule characteristics

4.1.1 The BGP Flowspec standard

The BGP Flowspec standard [5] specifies that a rule, i.e. flow specification, may consist of twelve components and one action. A packet matches a rule if it matches all com-

Type	Name	Example value	Description
1	Destination prefix	130.89.161.0/24	One CIDR prefix matching the destination address
2	Source prefix	130.89.161.0/24	One CIDR prefix matching the source address
3	IP protocol	1, 3, 5, 17-19	Any (range of) IP protocol numbers
4	Port	1-80, 443	Any (range of) TCP/UDP port numbers
5	Destination port	1-80, 443	Any (range of) TCP/UDP port numbers
6	Source port	1-80, 443	Any (range of) TCP/UDP port numbers
7	ICMP type	0, 3-5	Any (range of) ICMP types
8	ICMP code	3, 6-15	Any (range of) ICMP codes
9	TCP flags	ack, fin, push, syn	Any amount of TCP flags
10	Packet length	40, 255-1518	Any (range of) packet lengths in bytes
11	DSCP	40, 255-1518	Any (range of) DSCP bytes
12	IP fragmentation	dont-fragment, is-fragment	Any amount of IP fragmentation flags

Table 1. BGP Flowspec rule component types

ponents present in the rule. Table 1 shows all available BGP Flowspec rule components, including an example of their value and a description of the content they accept.

It is possible to provide four different actions, which would be applied to the packets matching the BGP Flowspec rule. If no action is specified, the IP traffic that matches a rule will be accepted. This research only considers the *traffic-rate* action, which allows rate limiting of IP traffic. Setting the rate limit to 0 will result in all traffic that matches the rule being discarded.

4.1.2 Rule definition

In this research paper, as well as in all programs written for this research, BGP Flowspec rules will be described as a dictionary, with the BGP Flowspec rule component type as the key with its corresponding value. The action of a rule will have its own nested dictionary, which contains the action type and its value.

The following example rule would discard all HTTP/HTTPS traffic to one of the Google web servers, coming from the *130.89.161.0/24* subnet:

```
{
  "type1": "172.217.19.195/32"
  "type2": "130.89.161.0/24"
  "type3": [6],
  "type5": [80, 443],
  "action": {
    "type": "traffic-rate",
    "value": "0"
  }
}
```

4.2 Limitations on BGP Flowspec rules

There are relevant practical limitations that should be taken into account when manually writing or automatically generating BGP Flowspec rules. These limitations stem from the BGP Flowspec standard itself, from the implementations of the standard and from the current routing hardware.

4.2.1 Limitations on the BGP Flowspec protocol

From an examination of the DDoSDB dataset, and from the definition of a DDoS attack, we can conclude that attacks often originate from many sources. We can see from table 1 that the BGP Flowspec standard allows specifying only a single source and destination prefix within a rule, while other rule component types allow for more flexibility by supporting multiple values. If one would want to block all sources in a DDoS attack using BGP Flowspec, that would result in an amount of rules equal to the amount of source IP addresses in the attack. This does not scale for large attacks and it is therefore needed to combine multiple sources into one prefix (section 4.3.2) or to ignore sources in the ruleset based on the amount of attack traffic they produce.

4.2.2 Limitations on hardware vendor implementations

Cisco, a major network hardware vendor, uses IOS XR as the operating system for their Series Aggregation Services platform [11], a hardware lineup designed for IXP-level routing. This operating system poses a limit of 3000 BGP Flowspec rules, along with a limit of five multi-value ranges within a BGP Flowspec Rule. Table 1 shows all rule component types that support multi-value ranges. Loibl et al. [8] show that the same hardware was unable to correctly dissect incoming BGP Flowspec rules larger than 239 bytes. It is unknown whether this bug has been fixed at the time of writing this paper.

Loibl et al. also make more general statements about the support and scalability of current routing hardware for the BGP Flowspec standard. They noted that "Internet routers are designed to keep a big destination based forwarding table (FIB) in their hardware, but when it comes to access control lists and forwarding policies, the underlaying hardware is much more limited and may not scale very well when a large number of flow specifications learnt from the entire Internet needs to be programmed." [8, p. 4]

Finally, all router vendors (Alcatel/Nokia, Juniper, Cisco and Huawei) that are used in the test lab of Loibl et al. [8] use a vendor-specific configuration language for BGP Flowspec rules. This makes automatic rule generation a more complicated process, as multiple output languages would have to be supported.

4.3 Implementation

An implementation has been created and published[1], which allows the use of the proposed system for generating rules and classifying impact. We will first explain our method of mapping DDoSDB fingerprints onto BGP Flowspec rules. Then, we propose methods for reducing the size of large rulesets, which is needed before they can be deployed in practice.

4.3.1 Mapping fingerprints onto BGP Flowspec rules

A direct mapping has been defined to generate BGP Flowspec rule candidates from the DDoS attack fingerprints. Table 2 shows each fingerprint attribute and the respective BGP Flowspec rule component it is mapped to. When multi-

[1] `https://github.com/DiedB/ResearchProject2019/blob/master/generator.py`

DDoSDB Fingerprint		BGP Flowspec rule	
Attribute	**Description**	**Type**	**Logic operator**
src_ips	List of source IP addresses	2	-
protocol	IP protocol	3	OR
dst_ports	List of destination ports	5	OR
src_ports	List of source ports	6	OR
additional.icmp_type	ICMP type	7	-
additional.tcp_flag	TCP flags	9	AND

Table 2. Mapping of DDoSDB fingerprint attributes onto BGP Flowspec rule component types

ple values are mapped into one rule component type, the logical operator that is used to separate these values is specified as well. All used rule component types support specification of multiple values separated by logical operators, except for the source IP address (type 2). This means that for every source IP address, a rule will be generated which contains that address and all other mapped fingerprint attributes.

Because source and destination addresses in BGP Flowspec rule can only be described as prefixes, the most basic implementation will map each source IP address onto a BGP Flowspec rule with a /32 prefix length. A more intelligent approach is described in section 4.3.2. DDoSDB anonymises the victim, and therefore no destination attributes can be mapped onto the BGP Flowspec rules. If destination data would be available, it would highly improve the quality of the generated rules.

Because a packet matches a rule only when all implemented component types match, it is desirable to add as much information about the attack into the rule as possible, given that this information is static during the attack. Examples of parameters that do not change during the attack are the protocol, the ICMP type(s) (if available) and the TCP flags (if available). This assumption holds because DDoSDB fingerprints only describe a single attack. Multi-vector attacks are dissected into multiple fingerprints, where each fingerprint describes a single attack vector.

For some DDoS attacks, the list of source ports or the list of destination ports can be static, meaning that each attacking packet is originating from the same source port or sent to the same destination port. For other attacks, either of these could be randomised and therefore have no value or even detrimental effects in a BGP Flowspec rule. Further investigation into the dataset shows that in the large majority of cases, fingerprints have either very few or a lot of source or destination ports. Therefore, ports are only mapped into a rule if a fingerprint when the total amount of source or destination ports is 5 or less.

The generated rules will always have a *traffic-rate* action with a rate of 0, effectively discarding all traffic that matches the rules.

4.3.2 Reducing ruleset size

After generating a ruleset using the method described in section 4.3.1, every source IP address that occurs in the attack has at least one corresponding BGP Flowspec rule. It follows from section 4.2.2 that this solution does not scale when attacks become larger in terms of source IP addresses. Therefore, a method is needed to reduce the size of the generated ruleset.

For type 1 and 2 rule components (destination prefix and source prefix), BGP Flowspec accepts addresses in CIDR prefix notation [12]. By decreasing the prefix length, it is possible to encapsulate multiple IP addresses in a single

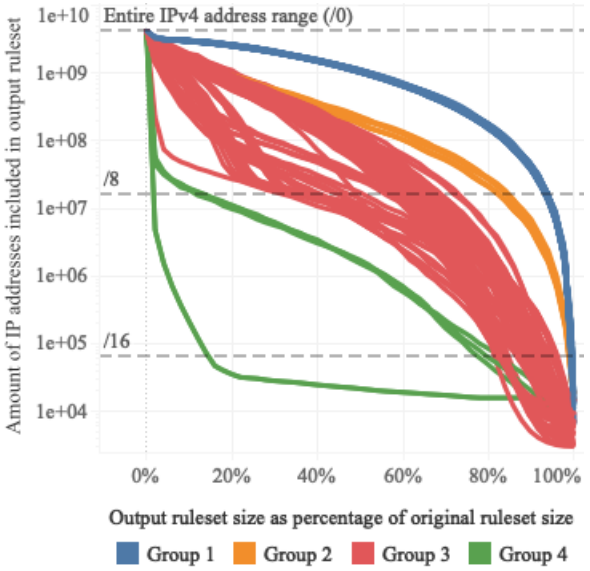

Figure 2. Reducing the ruleset size for large (>3000 source IP addresses) DDoS attack fingerprints

rule. In the large majority of cases, non-malicious IP addresses will be included in the prefix as well. This should not immediately cause a problem, since the associated impact is optimised during the quantification step and there are more matching criteria in rules to prevent all traffic from non-malicious IP addresses from being discarded. However, prevention of end-user impact is important and should be done wherever possible.

We present an algorithm that reduces the size of a list of IP addresses that is given as input, by bundling them into prefixes. In order to set a boundary to the size of the output list, the algorithm accepts an upper bound on the amount of output prefixes as input. The algorithm ensures that the total amount of IP addresses that are included in all output prefixes combined is always as low as possible.

Figure 2 shows the implications of reducing the input list of IP addresses into prefixes, for each of the 337 fingerprints in the DDoSDB dataset where the amount of source IP addresses is greater than 3000. This lower bound has been applied because the generated ruleset for these fingerprints will virtually always need reduction before it can be deployed in practice. The x-axis has been normalised to show the upper bound on the amount of output prefixes as a percentage of the original amount of IP addresses. The y-axis shows the amount of IP addresses that are included by all the prefixes in the output list combined. The trend lines show the amount of IP addresses covered by one prefix with the tagged length.

The resulting lines can be separated into 4 distinctive

groups, where group 3 (red) follows an average pattern and the other groups can be seen as outliers. These groups do not immediately share any properties, like a certain amount of source IP addresses or a certain protocol. However, we can conclude that the source IP addresses of group 1 (blue) and group 2 (orange) are relatively more distributed across the IPv4 address range than the source IP addresses in group 4 (green).

The best case that is represented in the chart is the innermost green line. It shows that an 80% reduction of the original IP address list is possible, without having many implications on the amount of IP addresses that is covered by the output prefixes. In turn, this means that a lot less rules are needed for this fingerprint while not having to increase the end-user impact by a large amount. The worst case is the outermost blue line, which shows that even a slight reduction means that a lot of IP addresses are covered by the resulting output prefixes, resulting in high end-user impact.

Altogether, this chart shows the potential of ruleset reduction over a set of large DDoS attack fingerprints, where some fingerprints allow for greater ruleset reduction without great implications on the resulting impact, while other fingerprints do not allow for much ruleset reduction. For these fingerprints, a different reduction method should be explored.

5. IMPACT QUANTIFICATION

Before deploying generated BGP Flowspec rules, their impact in the network needs to be quantified. This allows a network administrator to pick the most optimal rules, allowing for successful mitigation of a DDoS attack while minimising end-user impact. This section will explain the impact quantification process, by showing the factors that are used and how these factors are calculated and combined into a single quantification.

5.1 Factors for impact quantification

There are many factors that are relevant in quantifying the impact of a candidate BGP Flowspec rule. This section shows all impact factors that have been identified. We will explain why these factors are important, and how they can be computed in a real-world scenario.

5.1.1 Regular end-user impact

In most cases, the goal of a DDoS attack is bringing the target service offline. Failed mitigation efforts could block the DDoS attack in its entirety, while also reducing or denying service to normal users. In turn, this makes the DDoS attack reach its goal, even though it is not the attack traffic taking the service offline. This shows why it is important to keep track of regular end-user impact when using BGP Flowspec for DDoS attack mitigation.

5.1.2 Effectiveness

While preventing regular end-user impact is important, the effectiveness of a rule in discarding malicious traffic should also be assessed. A candidate BGP Flowspec rule could cause zero impact to non-malicious users, but it has no value when it does not block any malicious traffic.

5.1.3 Rule size

As shown in section 4.2.2, many hardware vendors pose limitations on the size and contents of BGP Flowspec rules. Even though a very advanced rule with lots of conditions can be very effective in mitigating DDoS attack traffic, it might become too large and complex to deploy to the network and therefore lose its value. Additionally, simpler rules are easier to read and comprehend by an ISP network administrator, allowing for manual modification when necessary.

5.2 Calculating the impact quantification

To enable fair comparison between multiple candidate BGP Flowspec rules, the factors for impact quantification that have been defined in section 5.1 have to be computed and weighted. Finally, the factors can be combined into one quantification by computing the weighted sum of all factors.

5.2.1 Calculating regular end-user impact

Regular end-user impact can be determined by simulating the deployment of the candidate rule on the traffic stream directed to the host that is experiencing the DDoS attack. It can then be calculated by measuring the amount of packets that are being discarded by the rule while not matching the characteristics of the attack that are described in its fingerprint, relative to the total amount of packets that match the described fingerprint.

This will result in a decimal fraction I which would be 0 when the BGP Flowspec rule does not block any normal traffic, and 1 for a BGP Flowspec rule that blocks all normal traffic.

5.2.2 Calculating effectiveness

The effectiveness of a BGP Flowspec rule can be determined by simulating the deployment of the candidate rule on the traffic stream directed to the host that is experiencing the DDoS attack. It can then be calculated by measuring the amount of traffic volume that is being discarded by the rule while also matching the characteristics of the attack that are described in its fingerprint, relative to the total traffic volume that matches the described fingerprint.

This will result in a decimal fraction E which would be 0 when the BGP Flowspec rule does not block any traffic, and 1 for a BGP Flowspec rule that completely mitigates the DDoS attack.

5.2.3 Calculating size

The size of a BGP Flowspec rule depends on the amount of rule component types it implements, and the amount of values associated to the implemented rule component types. It can be calculated using the following sum:

$$S_{abs} = \sum_{r \in R} f(v)$$

Where S_{abs} is the absolute size of the BGP Flowspec rule, R is the BGP Flowspec rule, v is the list of component values, and f is a function that returns the length of the list of values, counting each range as two values. Ranges should be counted as two values because of the underlying implementation of BGP Flowspec rules: a range between 1 and 100 would be defined as the operator-value pairs ≥ 1 and ≤ 100.

Before using the size in the final impact quantification, it needs to be expressed as a decimal fraction in terms of the largest allowed size using the following formula:

$$S = \frac{S_{abs}}{S_{max}}$$

Where S is the size of the BGP Flowspec rule expressed as a decimal fraction and S_{max} is the largest allowed size.

This will result in a decimal fraction which would be 0 for an empty BGP Flowspec rule and 1 for the largest rule

that is allowed in the rule generation process.

5.2.4 Calculating impact quantification

Once the individual impact factors for a rule have been computed, we can combine them into one quantification using a weighted sum.

The weights $w_{s,e,i}$ should be set based on the importance they have in the situation where the impact quantification is used to find optimal BGP Flowspec rules. It should be noted that the weights should be negative for factors that are detrimental to the quality of a BGP Flowspec rule, meaning that w_i and w_s should be negative.

Determining the weight of each factor and calculating the following weighted sum:

$$Q = w_s * S + w_e * E + w_i * I$$

Where Q is the impact quantification, S is the size, E is the effectiveness, I is the regular end-user impact and $w_{c,e,i}$ are the corresponding weights.

When calculating the impact quantification for a set of rules, the best rule would receive the highest impact quantification while the worst rule would receive the lowest impact quantification.

6. RESULTS

This section elaborates on the results of the combination of the presented solutions for BGP Flowspec rule generation and impact quantification. By taking example input data and simulating a real-world deployment of the solution, we can gather data about its performance and effectiveness while showing how a real-world deployment of the solution would work.

6.1 Solution overview

Figure 3 shows a schematic overview of the complete solution that has been presented in this paper. The combination of our contributions results in an integrated solution for ISP's that helps in using BGP Flowspec as a DDoS mitigation solution while preventing impact to regular network users as much as possible.

6.2 Real-world application

A simulation has been constructed which allows us to show how the presented contributions work together, and how they can be used in a real-world deployment of BGP Flowspec as a DDoS mitigation solution. A few usage scenarios will be presented, after which the setup of the simulation is explained. Finally, we will show how the presented solution of this paper would work for each of these usage scenarios.

6.2.1 Usage scenarios

ISP 1 is experiencing a large DDoS attack that cannot be handled by their core network infrastructure, resulting in a lot of collateral damage. There is no problem in the specific attacked web server going down, as long as the incoming traffic volume decreases significantly. This ISP only cares about rule effectivity while regular end-user impact is not as relevant.

ISP 2 is experiencing a DDoS attack on a critical server, where regular end-user impact should be avoided at all cost. Therefore, a lot more weight is on avoiding regular end-user impact rather than on the effectivity or size of the BGP Flowspec rules.

ISP 3 applies BGP Flowspec as a DDoS mitigation solution in combination with other solutions in their core network. They only want to block the sources that are producing the highest amount of volume and want to avoid end-user impact at all cost.

These scenarios yield the following weights:

ISP	w_s	w_e	w_i
1	-0.5	1	-0.1
2	-0.1	0.01	-1
3	-0.01	1	-1

6.2.2 Methodology

Since we do not have access to a live traffic feed of an ISP endpoint in an IXP network, we have to use alternative data sources as basis for our results. The regular end-user impact factor (section 5.2.1) uses live traffic data to determine the amount of **normal traffic** that is discarded by the BGP Flowspec rules, while the effectiveness factor (section 5.2.2) uses live traffic data to determine the amount of **attack traffic** that is discarded by the BGP Flowspec rules.

We utilise the TCPReplay dataset [10] to determine the amount of real-world traffic that would be matched by our generated BGP Flowspec rules. Since the TCPReplay dataset is assembled in a different network, the IP addresses in these packets do not match the IP addresses in our BGP Flowspec rules. Therefore, we remove the IP address information from our matching criteria.

To compensate for ignoring IP address data in traffic matching, we modify the regular end-user impact factor by multiplying it by a new factor. This factor represents the amount of IP addresses that are covered by the source prefix of the BGP Flowspec rule. This negatively impacts the score of BGP Flowspec rules that use larger prefix lengths, because these rules would theoretically match and discard more traffic. The modified regular end-user impact factor can be calculated using the following formula:

$$I_{new} = I * \frac{2^{32-p}}{2^{32}}$$

Where I_{new} is the impact factor used in this simulation, I is the impact factor that is calculated using the method described in section 5.2.1 and p is the prefix length used in the generated type 2 component of the BGP Flowspec rule.

The effectiveness factor E is also dependent on live traffic data. In the absence of this data, we use the trace file corresponding to the fingerprint provided by DDoSDB as traffic data.

The parameter r_{max} defines the maximum amount of rules that is allowed to be generated by the rule generation algorithm, and will be given a different value for each run of the simulation. r_max will be variated between 1 and 200 in the simulation.

6.2.3 DDoS attack data

As input for the presented solution, we use one DDoSDB attack fingerprint which describes a TCP SYN flood attack. A SYN flood works by starting a TCP handshake without ever finalising it, causing the server to keep a connection open. Given enough connection requests, the server's resources will eventually be depleted causing denial of service to normal users. The fingerprint contains 1721 source IP addresses and the corresponding trace file contains 622,650 packets. The attacked server is most probably a web server, given that all connection requests are sent to port 80.

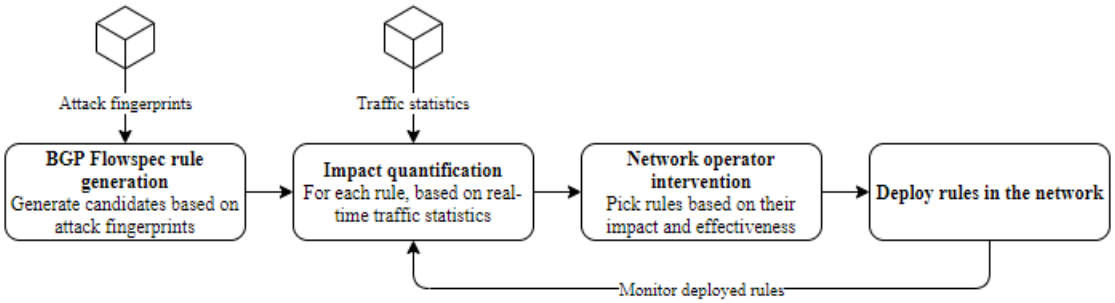

Figure 3. Overview of the proposed impact-based DDoS mitigation mechanism

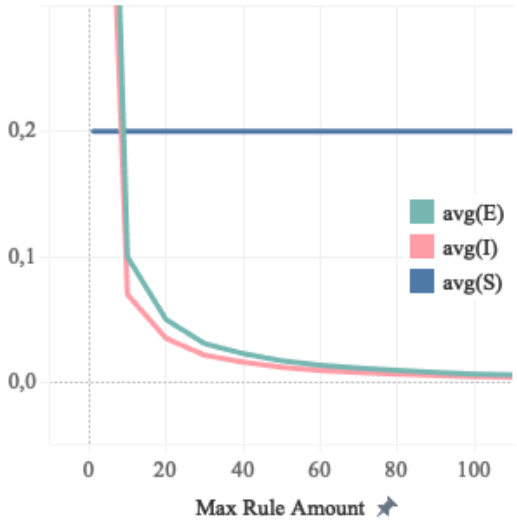

Figure 4. Effectiveness, size and general end-user impact of rules, while varying r_{max}

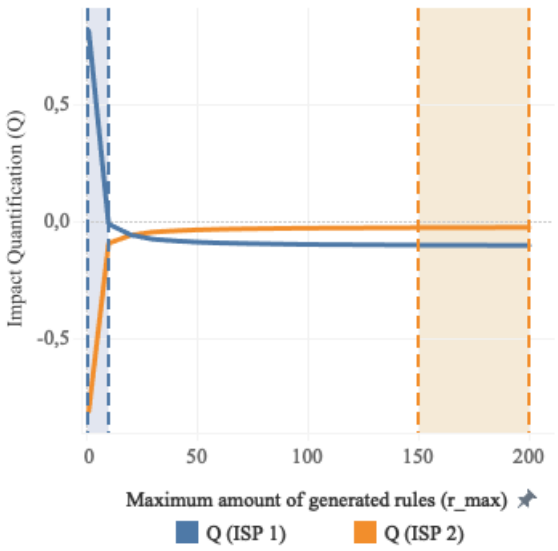

Figure 5. Impact quantification of rules, while varying r_{max}

6.2.4 *Characteristics of impact factors*

Figure 4 shows the average values of all impact factors when varying the size of the ruleset r_{max}. A few observations can be taken away from this:

1. While the effectiveness E of rules in a small ruleset is very high, this also results in high regular end-user impact I.
2. I will eventually reach 0, given a large enough ruleset
3. S is constant for every rule generated from this fingerprint, because the rule generation algorithm does not find any factors where variation on values is possible

In the case of this fingerprint, it would be beneficial to implement a larger ruleset if there is no constraint on the amount of BGP Flowspec rules that can be deployed into the network. However, because of the asymptotic behaviour of the curves for I and E, the returns of increasing the amount of deployed BGP Flowspec rules starts becoming negligible for larger rulesets. Another result that should be highlighted, is the fact that this fingerprint contains 1721 source IP addresses, but it is perfectly possible to succesfully mitigate the attack with fewer BGP Flowspec rules.

6.2.5 *Choosing candidate BGP Flowspec rules*

Figure 5 shows the average value of Q while varying the size of the ruleset, using the weights for ISP 1 and 2 as specified in section 6.2.1. It clearly shows that ISP 1 should be looking for candidate BGP Flowspec rules to deploy in the network within a ruleset size between 1 and 10, while ISP 2 should be looking at larger ruleset sizes to find an optimal mitigation approach.

ISP 3 should take a different approach in choosing rules, as this situation is mostly calling for outliers. These values, where E is relatively high while I is near zero, do not occur for this particular fingerprint.

7. CONCLUSION

In this research, we looked into methods of reducing impact to normal users as much as possible while deploying BGP Flowspec as a DDoS mitigation mechanism.

First, we presented the possibilities in BGP Flowspec rule definition and elaborated on the process of generating BGP Flowspec rules, using DDoS attack data from an existing platform for storing and sharing DDoS attack characteristics.

Then, we investigated methods of determining the positive and negative impact of BGP Flowspec rules when deployed in the network. We were able to identify important factors that constitute the amount of network impact of a BGP Flowspec rule, and presented a method of combining these into a single quantification.

Finally, we set up a simulation environment that allowed

demonstrating the presented solution. Although the choices that are made based on the impact quantification of a BGP Flowspec ruleset depend on a lot of situational factors, the added value of the presented solution has clearly been demonstrated.

All code that has been written in the process of this research, has been published as open source on GitHub[2].

8. FUTURE WORK

While an experiment has been conducted in which we test the practical implementation of the presented solution, it could only use a limited simulated environment where some assumptions had to be made. A real-world test should be conducted, with a focus on its performance and scalability.

The algorithm for rule generation and reduction that has been presented in this research, does not take any metadata about the input values into account. The quality of generated candidate rules would greatly improve if the algorithm could consider the amount of traffic volume originating from each source IP address. This way, more effective groups of IP addresses could be generated, by grouping IP addresses that generate large volumes of traffic and ignoring IP addresses that are not generating a lot of traffic.

Although BGP Flowspec shows lots of potential as a DDoS attack mitigation solution, it should in most cases not be used as the only line of defence against DDoS attacks. There are more stages, deeper into the ISP network, where DDoS attack traffic can successfully be discarded while applying more granular filtering. This does not mean that all mitigation efforts should be concentrated deeper into the network, as that would require a lot of computing power and could potentially cause more harm than good. More research could be done into using the presented BGP Flowspec-based DDoS mitigation solution in harmony with other DDoS mitigation efforts deeper into the ISP network.

9. ACKNOWLEDGEMENTS

We are very grateful to our supervisor Jair Santanna. He laid an important foundation for this research with the development of DDoSDB and his research into creating fingerprints for describing DDoS attacks. He also helped with various useful comments and discussions.

10. REFERENCES

[1] Oleg Kupreev, Ekaterina Badovskaya, and Alexander Gutnikov. *DDoS Attacks in Q3 2018*. `https://securelist.com/ddos-report-in-q3-2018/88617/`. Accessed: 2019-01-21.

[2] Cloudflare Inc. *Famous DDoS Attacks | The Largest DDoS Attacks Of All Time*. `https://www.cloudflare.com/learning/ddos/famous-ddos-attacks/`. Accessed: 2019-01-21.

[3] José Jair Santanna et al. "Booters - An analysis of DDoS-as-a-service attacks". In: *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on*. IEEE. 2015, pp. 243–251.

[4] Jessica Steinberger et al. *LUDO–KIDS PLAYING DISTRIBUTED DENIAL OF SERVICE*.

[5] Pedro Marques et al. *Dissemination of flow specification rules*. RFC 5575. RFC Editor, 2009. URL: `http://www.rfc-editor.org/rfc/rfc5575.txt`.

[6] Nico Hinze et al. "On the Potential of BGP Flowspec for DDoS Mitigation at Two Sources: ISP and IXP". In: *Proceedings of the ACM SIGCOMM 2018 Conference on Posters and Demos*. ACM. 2018, pp. 57–59.

[7] Lennart van Gijtenbeek and Tim Dijkhuizen. *DDoS Defense Mechanisms for IXP Infrastructures*.

[8] Christoph Loibl and Martin Bacher. *BGP Flow Specification Multi Vendor and Inter AS Interoperability*.

[9] Jessica Steinberger et al. "Collaborative attack mitigation and response: a survey". In: *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on*. IEEE, pp. 910–913.

[10] Fred Klassen and Appneta. *TCPReplay - Sample Captures*. `http://tcpreplay.appneta.com/wiki/captures.html`. Accessed: 2019-01-27.

[11] Inc. Cisco Systems. *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide, Release 5.2.x | Implementing BGP Flowspec*. `https://bit.ly/2FVTUgF`. Accessed: 2019-01-22.

[12] T. Li V. Fuller. *Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan*. RFC 4632. RFC Editor, Aug. 2006. URL: `http://www.rfc-editor.org/rfc/rfc4632.txt`.

[2] `https://github.com/DiedB/ResearchProject2019`