

# An Adaptive Anti-DDoS System for IP Service Provider Backbones

Alberto Savio Figueira Rodrigues<sup>1</sup>, Fabio Santos<sup>1</sup>, Marcos Araujo<sup>1</sup>,  
Natalia Castro Fernandes<sup>1</sup>

<sup>1</sup>Grupo Midiacom - PPGEET/Escola de Engenharia  
Universidade Federal Fluminense (UFF)  
Niterói – RJ – Brasil

{savio, santosfabio, marcosaraujo, nataliacf}@id.uff.br

**Abstract.** *This article presents an adaptive anti-DDoS system based on SDN for IP service providers' backbones. We analyzed requirements and solutions used on current IP backbones in order to manage different anti-DDoS systems to mitigate attacks to customers or to the backbone itself. The base of our proposal is a controller that synchronizes network reaction according to the volume of attacks and the available infrastructure, using a layered protection scheme. This controller dynamically provisions virtual machines and network links based on transit virtual router and forwarding (VRF). The system is able to dynamically reconfigure itself according to attack traffic patterns. As a consequence, our system improves backbone performance and customers quality of experience by reducing the impact of DDoS traffic in a more efficient way than current solutions.*

## 1. Introduction

All networks connected to Internet are exposed to attacks. Service Providers Network Operators have to start actions against Denial of Service (DoS) and Distributed DoS (DDoS) attacks, targeted to the backbone itself and to clients. Implementing protection systems at network edges isn't enough, since botnets and volumetric attacks can damage the network even by overloading servers processing power or links [Ferguson and Senie 2000]. As a consequence, service providers (SP) must get involved on countermeasures, using both traditional solutions and improved ones. Countermeasures to volumetric attacks should be activated in layers, from the customer firewall to customer link, inside SP backbone and at the backbone border. Sometimes simple agreements between SPs can extend drop measures beyond the SP backbone borders.

Monitoring, configuring and adjusting the network and systems used to detect and react to attacks manually can be ineffective, require a lot of effort and increase costs. SP operators also need to know that an attack is on the run regardless of the volume of traffic. After that, the operators define the best traffic filtering point and create all the required infrastructure to activate the traffic verification.

This article analyzes the main solutions currently being applied against DDoS on SP backbones, highlighting the weaknesses of these solutions in stand-alone mode. Based on this analysis, we propose to integrate solutions with interaction between network and anti-DDoS systems under a software defined network (SDN) view. we understand an

anti-ddos system as a set of softwares and hardware dedicated to the mitigation of DOS and DDOS attacks. Our solution works in a layered model, controlled by a centralized system.

In our proposal, a central controller dynamically activates virtual routing and forwarding (VRFs) [Rosen and Rekhter 2006] to reroute traffic flows to different levels of anti-DDoS verification centers. In order to deal with dynamic traffic volume, our controller creates instances of virtual machines (VMs) to evaluate the traffic, filter the unwanted flows, and return the legitimate traffic to the destination on the network. Whenever a DDoS traffic is detected, the system promotes a self-reconfiguration to attend the new demand and to guarantee that the verification points are as close as possible to attack flow source in the backbone. Besides controlling VRFs, our controller monitors VM usage and traffic volume in order to autonomously create and destroy VMs and to properly distribute the traffic flows among servers. Moreover, when the initial system, which is based on full-time packet inspection is overloaded, the system starts different anti-DDoS systems in order to provide a fast and efficient response against the attacks in the backbone. We perform a qualitative analysis comparing our system to other proposals that are currently used in service provider networks.

The remaining of this paper is organized as follows. Section 2 presents the related work. Section 3 presents the main schemes currently used to protect backbone networks against DDoS attacks. Section 4 describes the proposed system and Section 5 presents a qualitative analysis. Finally Section 6 concludes the paper.

## **2. Related Work**

Distributed Denial of Service (DDoS) attacks are one of the biggest concerns for security professionals [Jonker et al. 2016, Gillman et al. 2015, Zargar et al. 2013]. Attackers can turn services and even interactive networks unavailable to users using tools like botnets. Nevertheless the DDoS attack traffic is not always perceptible to network operators, due to the distributed nature of these attacks.

Ferguson and Senie discuss a method for using ingress traffic filtering to stop DoS attacks that use forged IP addresses to be propagated from 'behind' an Internet Service Provider's (ISP) aggregation point [Ferguson and Senie 2000]. Some authors discuss the impact of DDoS attacks in cloud computing environments [Santanna et al. 2015, Yan et al. 2016]. With recent advances in software-defined networking (SDN), SDN-based cloud brings new chances to defeat DDoS attacks in cloud computing environments.

Jakaria et al. proposes the use of Network Function Virtualization (NFV) technology to perform balance with the attack load. The system creates dynamic network agents to intercept packets when the system experiences DDoS attack. The proposed algorithm defines the number of agents according to the attack traffic [Jakaria et al. 2017].

In this way our proposal is directly related to using distributed and automatic structures to detect DDoS attacks based on data flow monitoring by a central controller. In our proposal, instead of discussing how to detect the attack, we propose an efficient method to guarantee that the malicious traffic is filtered in the verification center near to attack source, under the point of view of the backbone, and the legitimate traffic is correctly routed to the destination. To this end, our proposal is based on automation and virtualization technologies, without restrictions concerning the scenario like

in [Santanna et al. 2015, Yan et al. 2016]. Our proposal not only defines the number of virtual machines used to inspect packets, as proposed in [Jakaria et al. 2017], but also their position in the network. Besides it also deals with traffic redirection in an autonomous way, according to the position of the chosen verification center.

### **3. Current DDoS prevention methods**

Nowadays security measures for backbone networks apply to routing protocols and bandwidth configuration. Blackholing is a simple technique where some pre-defined routers announce themselves to be the next-hop to an attacked IP address/range. Then the routers drop the received traffic. Indeed, the blackholing technique was expanded to allow a distributed action (RFC3882), where instead of sending undesired traffic to a specific router (BLK router) to be dropped, the dropping action can be performed at any backbone router.

Some of the solutions allows customers to activate themselves the blackholing. This prevents customer traffic marked as malicious or undesired to traverse the network and get to the customer port. Although effective against DDoS, the use of blackholing is very restrictive, as it does not try to differentiate legitimate and DDoS traffic, but drops all the packets to a destination. Hence, more sophisticated solutions than BLK are required to protect backbone networks.

#### **3.1. NetFlow and traffic inspection**

Manually activating traffic discard via blackholing is too slow. Then NetFlow, defined on RFC 3954, collects traffic samples and sends them to inspecting tools. If an attack is detected, the verification system can automatically send to the network a poison route that will remotely start a distributed blackholing. Although this can be very effective, it requires that the operator sets NetFlow collecting points at selected devices. Generally these points are peering connections to other SPs Backbones and the router where the customer is connected, in some special cases.

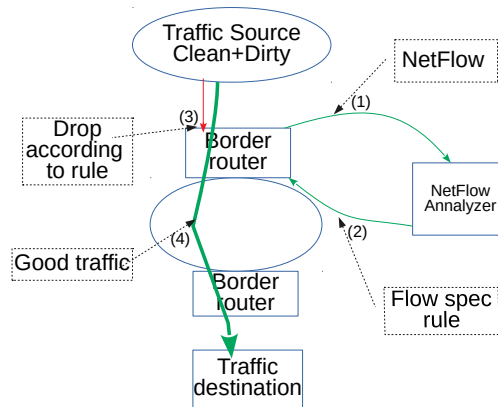
The issue with this solution, based on blackholing, is that all traffic to a specific customer IP range, malicious or not, is discarded. Hence, the network management creates a case of DoS, at least for this IP range.

#### **3.2. BGP FlowSPEC and Cleaning Centers**

BGP FlowSpec, defined on RFC 5575, is a feature that allows a more selective discard at the border of the network. It's possible to activate this service manually or via a traffic inspection tool that receives NetFlow data. When a threat is detected, BGP FlowSPEC policies are generated and sent to routers, via direct BGP sessions or via Route reflectors. The available features depend on the vendor implementation, but usually routers can match packet headers on layer 3 or higher, source or destination.

Figure 1 shows the basic procedures of BGP FlowSPEC. The border router sends traffic samples via Netflow to the NetFlow Analyzer (1). If an attack is detected, a FlowSpec rule is sent via BGP to the router (2). Hence, the undesired traffic is dropped or shaped (3), and good traffic is sent across network to its destination (4).

Instead of dropping packets, another option is to forward the traffic to a deep inspection tool. In this case, the Netflow analyzer, after detecting a menace, sends a



**Figura 1. BGP FlowSPEC overview.**

“poison route”, that contain the attacked destination IP range and changes the next-hop to a “cleaning center”. This structure analyzes and separates legitimate traffic from the malicious packets or flows. After that the legitimate traffic must be sent back to network in a way that it is delivered to the real destination.

The first solution to return the traffic to the network is to use tunneling. The cleaning center starts a tunnel (generally some IP over IP) from the Cleaning Center to the customer device. As the network only sees the external packets, the poison route does not send the packets back to the cleaning center.

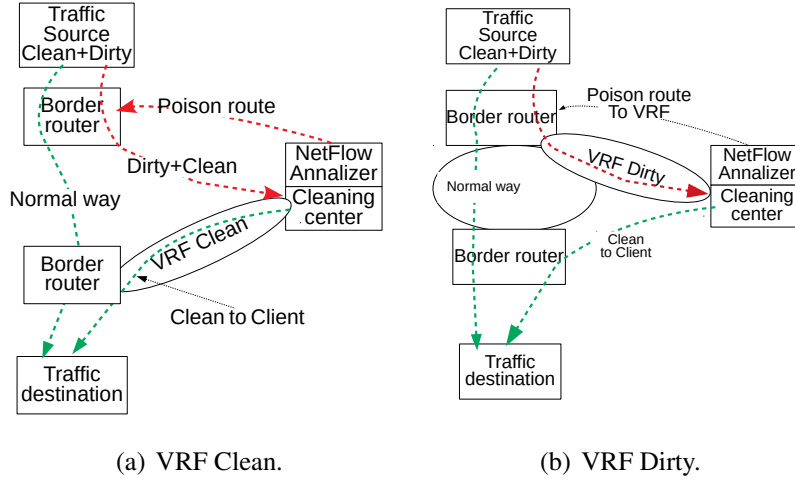
Bi-directional solutions can be implemented via more than one backbone. However it demands many resources on the cleaning centers servers to establish and forward traffic over several tunnels. If the tunnel connection check is disabled, this solution becomes lighter, but it exposes the customer network. Hence, some customers refuse tunnel-like solutions to avoid vulnerabilities in their networks.

### 3.3. VRF CLEAN versus DIRTY and cleaning centers

For most multi service networks that implement Multi-protocol Label Switching (MPLS) and Multi-Protocol BGP (MPBGP), MPLS Virtual Private Networks (MPLS-VPN) are common structures. In these cases, the idea is to use MPLS-VPN VRFs to transport CLEAN traffic (classified as legitimate) or DIRTY traffic (classified as malicious) from the inspection tool back to the network, still keeping the strategy of the poison routes sent by the NetFlow analyzers. Clean VRFs are used to send back the cleaned traffic to customer, as shown on Figure 2(a). VRF Dirty, shown on Figure 2(b), is the opposite of VRF clean. In VRF Dirty, the poison route has as a next-hop (cleaning center) that is inside the VRF Dirty. After been cleaned, the traffic is sent back to the global routing domain.

### 3.4. Virtualization and full inspection

There are some disadvantages when using Netflow to monitor attacks: NetFlow operates under statistical packet capture and there is a delay before the poison route can be injected and computed on routing protocols. A continuous traffic analyzer associated to dropping police is usually much more efficient. To do that an operator could install a dedicated



**Figure 2. Overview of the usage of VRF.**

device at the customer port on the backbone, but this would be expensive. Instead using a shared platform that deals with several clients' traffic is an affordable solution. Hence the operator uses a virtualization system to activate a virtual traffic analyzer per customer. As attacks do not happen at the same time for all clients, it is possible to expand and shrink instance capacity according to traffic volume. In this scheme customer traffic is redirected to a virtual traffic analyzer and after inspected and cleaned, it is sent back to customer. Again, it is possible to use both tunneling or VPN. For this kind of service, each customer has a separated VRF.

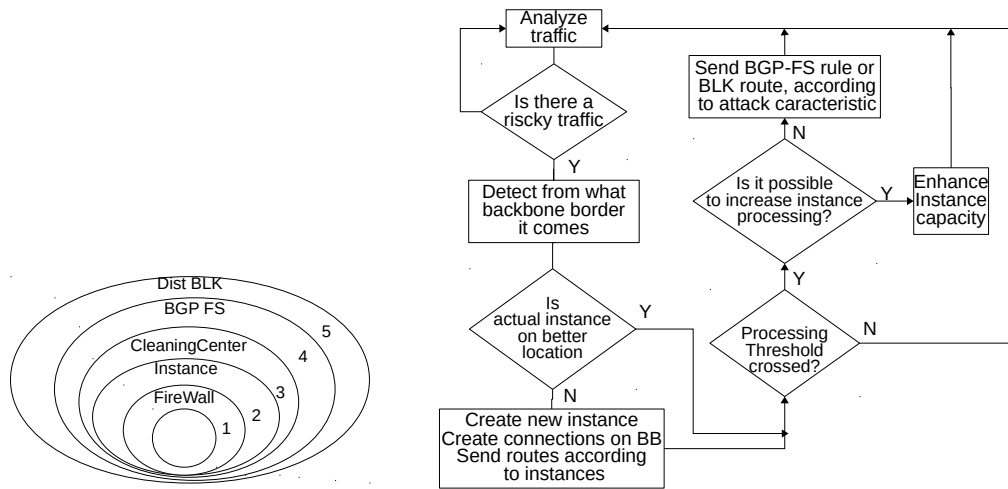
#### 4. Proposal Description

As we discussed, there are different anti-DDoS techniques for backbone networks. Table 1 summarizes all the described techniques. Nevertheless, these methods are not integrated and have a static instance placing for traffic analysis systems. Besides, the analysis systems usually present a fixed capacity, even though the processing requirements, which depends on attack traffic patterns, vary with time.

Method	Actions
Blackholing	Drop all packets for a destination
Distributed Blackholing	Drop all packets for a destination
BGP FlowSpec	Drop packets, source or destination match
Cleaning Centers	Drops dangerous traffic
Virtualization and full inspection	Deep analysis and drop

**Tabela 1. Current anti-DDoS methods used in backbone networks.**

In this scenario, we propose to use a cloud-based platform to activate virtual traffic analyzer instances on sites closer to customer. This is important on countries like Brazil, due to augmented delays caused by distance between customer routers and traffic analyzers, which reduces TCP protocol performance. With a good controller, it would be possible to have more than one virtual device per customer, not only for redundancy but to spread processing demand.



(a) Integrated protection sequence.

(b) Multi layers action algorithm.

**Figure 3. Proposed mechanism for adaptive protection against DDoS attacks in backbone networks.**

We developed an algorithm to control virtual traffic analyzer location for each customer, according to attack traffic origin and volume, on an adaptive way. The virtual traffic analyzer is a software running on a virtual machine that executes full-time packet inspection. Our system observes traffic changes, automatically adjusting the location of the virtual traffic analyzers according to customer IP ranges advertisements. Also, based on traffic volume, our system provisions virtual machine instances, adjusting the number of virtual machines, as well as their processing and memory capacities.

Our system is based on a multi-layer action algorithm. We use an integrated protection sequence, as shown in Figure 3(a). This sequence guarantees a scalable protection system, considering that the protection scheme should vary according to the number of attacks and the remaining processing power for detecting attacks. Specifically, the firewall is the first protection barrier that should be applied to all flows. Next, the accepted flows are directed to the closest virtual traffic analyzer, which performs a full packet inspection and decides which flows should be forwarded or dropped. In case there are a huge number of simultaneous attacks and the virtual machine servers are overloaded, part of the traffic is redirected to the cleaning center. Redirecting the traffic to the cleaning center consumes bandwidth resources of the backbone and should be avoided. Nevertheless it is the most proper solution in such cases, as the cleaning center is developed on real machines, which have more processing power than the virtual traffic analyzers. In case the cleaning center detects that some flows should be dropped instead of filtered, our system activates the BGP FlowSpec solution or the distributed blackholing, depending on the granularity of the flow that should be discarded. Hence, the external layers of our integrated protection system automatically causes the attack traffic to be discarded on the edge routers, which also reduces the load on the internal layers.

Figure 3(b) presents the flow chart of the proposed algorithm. The first step is to perform the traffic analysis using the virtual traffic analyzer. Whenever a new DoS or

DDoS flow is found, besides dropping the flow, the proposed system reorganizes itself based on attack traffic features and on the available resources for traffic analysis and dropping. When the reorganization process is triggered, our system first detects the origin in the backbone of the DoS or DDoS traffic, specifying a router or a set of routers. Then the system evaluates whether the virtual traffic analyzer instance that identified the flow is on the best location on the network. In case a new instance can be deployed closer to the router receiving the malicious traffic, the system automatically provisions the new instance, the new connections in the backbone and the corresponding routing announces. Next, the system analyzes virtual machine (VM) resource usage. The system defines processing thresholds for each virtual traffic analyzer. In case the number of flows directed to the VM containing the virtual traffic analyzer causes the processing rates to be greater the defined threshold, the system tries to dynamically change the amount of resources on the VM. In case there are no more physical resources available in the server to increase VM processing power, the system evaluate the nature of the DoS flow, deciding between a more granular packet discard using BGP FlowSpec or a more aggressive action using blackholing.

## 5. Analysis

We performed a qualitative analysis of the proposed system. We considered the pros and cons of each anti-DDoS solution for backbone networks and compare then to our system features. Table 2 shows the results.

Method	Pros	Cons
BlackHoling	Extremely simple.	Drop all; cross all network.
Distributed Blackholing	Simple; Backbone border.	Drops all.
BGP FlowSpec	More precise then BLK; Backbone border.	Need flow collector and analysis system; Down stream based.
Cleaning Centers	More precise then BGP-FS; Forwards traffic only on attacks.	Need flow collector and analysis System; Down stream based; More expensive; Needs extra hardware.
Virtualized Full Inspection	Symmetric or asymmetric; Full time inspection; No NetFlow collector; Shared hardware.	Full time traffic forwarding(cost); All clients' traffic forwarded to the same place;
Proposed System	Symmetric or asymmetric; Full time inspection; No NetFlow collector; Shared hardware; Better instance location; Dynamic capacity per instance.	Full time traffic forwarding(cost); Controller complexity.

**Tabela 2. Comparison of anti-DDoS methods for backbone networks.**

We observe that our proposal brings new advantages when compared to the other techniques. Indeed the proposed algorithms allows a better choice for virtual traffic analyzer instance, which is located as close as possible to the edge routes. It allows a full-time

traffic inspection, which results on more accurate results. Our proposal is also more complex than the previous methods, since it requires an intelligent controller which is able to choose between the available DoS filtering/dropping technique, it is able to control the network using SDN, it creates and destroys virtual machine instances at servers in different locations.

## 6. Conclusion & Future Works

The Internet is rapidly growing, increasing the number of hosts and the diversity of services in the network. The number of attacks, specially DDoS attacks, is also increasing. In this scenario, even a smart phone can become part of a botnet.

We proposed a system that provides a really effective protection against DDoS in SP networks. Our proposal includes layers of protection whose action is activated according to the attack level. If a single device can't deal with the attack, the system spreads the protection over the network, using cloud resources, BGP FlowSpec and blackholing.

Our system is currently under development. We have an initial prototype which is under test using a network composed of devices of different vendors. We are using SDN based on Netconf & Yang to create VRFs and OpenStack to manage virtualized systems.

## Referências

- Ferguson, P. and Senie, D. (2000). Network ingress filtering: Defeating denial of service attacks which employ ip source address spoofing. RFC 2827.
- Gillman, D., Lin, Y., Maggs, B., and Sitaraman, R. K. (2015). Protecting websites from attack with secure delivery networks. *Computer*, 48(4):26–34.
- Jakaria, A., Rashidi, B., Rahman, M. A., Fung, C., and Yang, W. (2017). Dynamic ddos defense resource allocation using network function virtualization. In *Proceedings of the ACM International Workshop on Security in Software Defined Networks & Network Function Virtualization, SDN-NFVSec'17*, pages 37–42, New York, NY, USA. ACM.
- Jonker, M., Sperotto, A., van Rijswijk-Deij, R., Sadre, R., and Pras, A. (2016). Measuring the adoption of ddos protection services. In *Proceedings of the 2016 Internet Measurement Conference, IMC '16*, pages 279–285, New York, NY, USA. ACM.
- Rosen, E. and Rekhter, Y. (2006). BGP/MPLS IP Virtual Private Networks (VPNs). RFC 4364.
- Santanna, J. J., van Rijswijk-Deij, R., Hofstede, R., Sperotto, A., Wierbosch, M., Granville, L. Z., and Pras, A. (2015). Booters — An analysis of DDoS-as-a-service attacks. In *2015 IFIP/IEEE International Symposium on Integrated Network Management (IM)*, pages 243–251.
- Yan, Q., Yu, F. R., Gong, Q., and Li, J. (2016). Software-Defined Networking (SDN) and Distributed Denial of Service (DDoS) Attacks in Cloud Computing Environments: A Survey, Some Research Issues, and Challenges. *IEEE Communications Surveys Tutorials*, 18(1):602–622.
- Zargar, S. T., Joshi, J., and Tipper, D. (2013). A Survey of Defense Mechanisms Against Distributed Denial of Service (DDoS) Flooding Attacks. *IEEE Communications Surveys Tutorials*, 15(4):2046–2069.