allot
See. Control. Secure.

DDoS and NFV are a Perfect Match

Position Paper

April 2020

See. Control. Secure.

Contents

1 Introduction

DDoS (Distributed Denial of Service) is here and it's not going away! It seems that every month we hear about a new attack, and it's not surprising that many types of DDoS attacks are referred to as floods —there is even one called a Tsunami — because their impact is overwhelming. They inundate network resources, including elements such as firewalls that are intended to ensure network security. Everyone with a web presence or internet connectivity is a potential target.

From quarter to quarter, attack profiles change. Kaspersky reported that while the number of attacks fell in Q419, their duration rose significantly compared to the previous quarter.

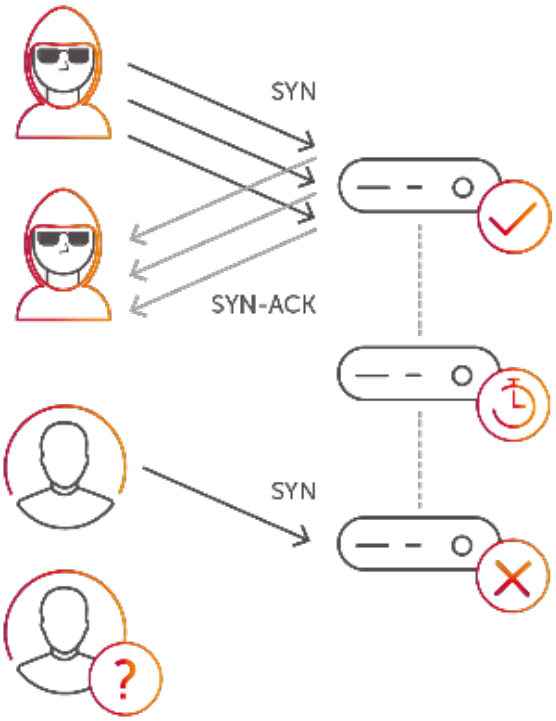

Large-scale volumetric attacks have also received extensive media coverage as they threaten online services of every kind. Even when communication service providers (CSPs) are not the intended target, their services suffer as the attack traffic passes through their network on their way to the actual victim. Extreme traffic volume can impede service delivery and even crash switches, routers, servers, and other network elements. The costs are high to CSPs and to their consumer and enterprise customers.

There are two dominant approaches to DDoS detection and mitigation: scrubbing centers and inline solutions. Each has its pros and cons, which are discussed later in this paper. In general, inline solutions provide faster, more thorough mitigation. However, they can be costly as an instance must be deployed at every peering point for complete coverage.

The good news is that NFV shakes up this paradigm. Virtualization and edge computing, as will be explained below, significantly lower deployment costs by enabling cost-effective, demand-based utilization of virtualized resources.

2 DDoS is a multi-faceted problem

2.1 Different attack directions

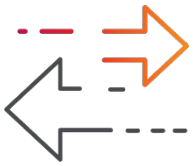

It is important to note that when people talk and think about DDoS they are generally thinking about incoming attacks where they are the targets. There is another kind of DDoS attack that CSPs must consider as well. These are the outgoing attacks where home consumer devices, enterprise/IoT devices, and mobile devices of all kinds, are subverted to launch attacks via your CSP network, making you, the CSP, appear to be the source of the attack. This can seriously harm your reputation and can potentially bring down your network services.

2.2 Reason for growth of the phenomenon

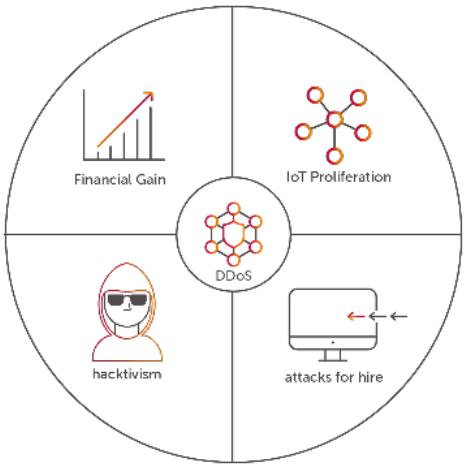

Why are DDoS attacks growing both in size and in frequency? A key reason is easy financial gain. DDoS attacks may include ransom demands, or may simply harm a competitor, either by crippling their business, or as a smokescreen to hide theft of business secrets. Another reason is that the explosive proliferation of IoT devices provides hackers with a growing landscape from which to launch attacks. Most IoT devices are easily hacked and converted into botnets, triggering ever-growing floods of DDoS assaults. A third reason is that hacktivists see DDoS attacks as an easy way to punish ideological enemies, be they government or corporate. DDoS attacks can also be a form of nation-state cyberwarfare, used both to harm operational capabilities and as a smokescreen to hide a subsequent theft of state secrets.

Finally, although DDoS is technologically sophisticated, it is surprisingly easy to launch an attack. There is a huge industry of ready-to-use DDoS attack botnets for hire, costing as little as 100 dollars per attack.

2.3 Impact on business

DDoS attacks are more than an inconvenience; they cause huge, and growing, financial costs. According to various industry sources, it costs over $120K to recover from small attacks. Recovering from large attacks often costs between $1M - $2M and sometimes more. Costs are, of course, both direct and indirect. They can include SLA penalties, costs incurred by overloaded call centers, and costs related to restoring/replacing

damaged infrastructure. On top of these, there may be significant costs associated with repairing damage to reputation: customer churn, promotional campaigns, and discounts offered to win back customers, or to attract new customers.

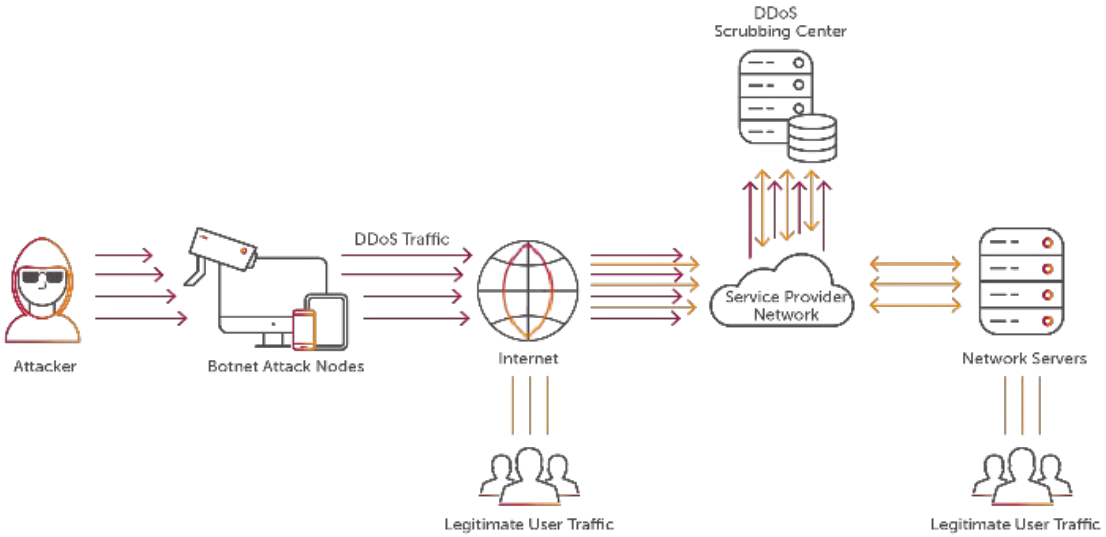

3 Pros and cons of scrubbing center and inline solutions

The two main approaches to mitigating DDoS traffic are scrubbing center and inline solutions. Scrubbing center solutions detect attacks and reroute all of the traffic to specialized infrastructure where the attack traffic is removed, and the clean traffic is routed back into the CSP network. Inline solutions detect and drop DDoS traffic within the CSP links, allowing only clean traffic to continue on its journey. Each approach has its advantages and disadvantages.

3.1 Scrubbing center pros and cons

Scrubbing center solutions are usually deployed and maintained by third party vendors, thus the upfront and ongoing expenses of deploying and running the solution belong to someone else. It is generally inefficient to route all of the traffic to scrubbing centers, so Cisco NetFlow, or an equivalent process, is used to sample the traffic and, when attacks are detected, the traffic is routed to the scrubbing center. The hardware footprint of the NetFlow probes can be smaller and less expensive than a solution that inspects all traffic. There can be significant CAPEX outlay for probes and any extra routing and transport infrastructure costs, but the ongoing cost to the CSP is largely OPEX, which is usually easier to accommodate in today's business environment.

The disadvantages of scrubbing center solutions are directly related to their remoteness:

- **Lowered Performance** – the extra data journey to and from the scrubbing center introduces increased latency, jitter, and packet loss to the legitimate traffic. This degrades the user experience, especially in performance-sensitive applications like VoIP and streaming video.
- **Delayed Implementation** – rerouting requires network routers to publish and propagate new routes (BGP/OSPF, etc.) in order to redirect all traffic to the scrubbing center. This can take 2-3 minutes, which could cause serious damage during a large attack. For many new-wave "hit and run" attacks, this level of delay is unacceptable.
- **Less Comprehensive** – because NetFlow-based solutions only sample the traffic, they cannot provide 100% effective attack detection. Similarly, they are not able to detect low-rate, application-based attacks. Finally, scrubbing center solutions are typically implemented to inspect incoming traffic, thereby missing the growing incidents of outbound attacks.

3.2 Inline pros and cons

Inline DDoS protection solutions enable attack detection and surgical mitigation on the spot, without regard for size or duration of attack and without diverting huge volumes of legitimate traffic and introducing delays. Because these solutions are inline, they can inspect outbound traffic as well as inbound traffic, enabling correlation of traffic flows, which improves accuracy and reduces the incidence of false positive or negative identification. In addition, they prevent CSPs from being a conduit for attacks originating on devices connected to their network.

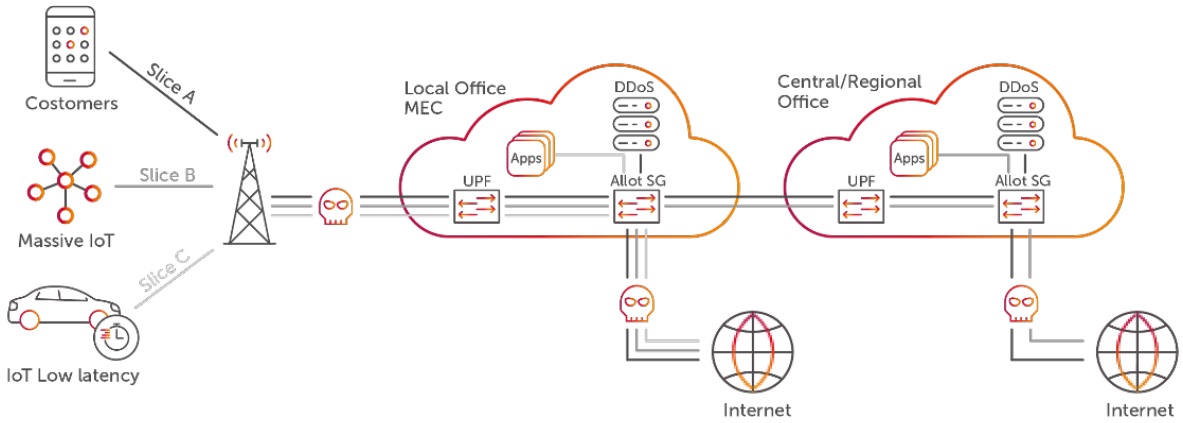

The principal, significant disadvantage of inline solutions is that, because they monitor all traffic and perform mitigation at the point of detection, they require carrier-grade capacity, throughput, reliability, and scalability and must be deployed at every peering point to keep malicious traffic out of the CSP network. Therefore, it requires a bigger up-front capital expense to deploy throughout the CSP network than is required by scrubbing center solutions, especially cloud-based scrubbers.

4 How can NFV help with DDoS mitigation?

In a nutshell, NFV (network function virtualization) is meant to lower CAPEX by avoiding over-allocation of dedicated hardware to meet every function's worst-case scenario, and to lower OPEX by enabling centralized deployment and management of services and remote troubleshooting of network function problems.

5 Impact of NFV on the limitations of inline detection & mitigation

Even without 5G, DDoS attacks have grown to be enormous, such as the 2018 memcached attack. Lately, they have been smaller, taking a more hit-and-run approach to try and avoid mitigation solutions that work by sampling the traffic. But the coming exponential spread of high-speed bandwidth promised by 5G means that dramatically bigger attacks will be possible because the "5G highway" will have many more lanes to enable vastly higher rates of traffic, both good and bad.

If inline DDoS detection and mitigation solutions are costly in 4G, they could be cost-prohibitive in 5G, due to the expected growth in data traffic and the resulting capacity demands of worst-case provisioning. In the case of DDoS protection, CSPs would need to provision worst-case compute resources to handle the largest possible volumetric attack at every peering point. This would certainly be economically unfeasible. However, the concepts of multi-access edge computing (MEC) and NFV can be combined to enable a cost-effective solution.

If inline DDoS detection and mitigation are implemented via the MEC paradigm, the attacks can be mitigated as close to the source as possible, keeping harmful traffic from getting past the edge into the core of the CSP network, thereby minimizing its impact. However, this is still a very costly solution, if worst-case solutions must be deployed at the edge for every peering point of the network.

This is where NFV comes into play. By virtualizing the DDoS mitigation network function, it becomes possible to under-provision the worst-case scenario defense. NFV makes it possible to scale out the DDoS solution, at the exact edge location and at the exact scale needed to meet and mitigate even the worst attack. Because multiple edge compute functions will all be deployed using NFV, their shared spare resources are less than if each function needed to pre-allocate its full worst-case set of resources. Thus, the inline DDoS detection and mitigation solution becomes much less expensive as it utilizes shared, excess resources among a host of edge compute functions that are all virtualized.