Pipeline
Technology for Service Providers.
www.pipelinepub.com Volume 14, Issue 2

How NFV Can Enable 'Pay-as-you-Protect' Network Protection

By: Nicolas St. Pierre

The Challenge

The scale of malicious attacks against communications service providers (CSPs) is increasing at a frightening pace. According to a report from Arbor Networks report, in the first half of 2016, the peak DDoS attack size reached 579Gbps, up 73 percent from 2015. There were 46 attacks over 200Gbps monitored in the first half of 2016, versus 16 during all of 2015. Meanwhile, despite massive growth in attack size at the top end, 80 percent of all attacks are still less than 1Gbps.

Tier-1 CSPs understand first-hand that legacy network protection solutions, such as those for DDoS attacks, create enormous inefficiencies, as they require proprietary hardware that is scaled for the rare peaks (579Gbps) and that otherwise either sit idle or are massively overprovisioned as they battle much smaller attacks (1Gbps). Even then, a CSP cannot be sure that the scale will be sufficient when the next attack peak comes.

The costs associated with such large-scale attack mitigation (or “scrubbing”) platforms — including those related to hardware, network resources, IT personnel, and routing infrastructure — are an incredible burden for any network operator, particularly Tier-1 CSPs, given their size.

The answer to this growing problem may lie in virtual network functions (VNFs) that can automatically scale up only when needed, and to the exact size needed, then be broken down when no longer needed — true elastic scaling. Resources are paid for only when used and to the extent used. While this concept may look promising on a whiteboard, I was challenged by an Asian tier-1 CSP to design a proof-of-concept that would demonstrate the viability of such an approach. So I did.

The Evolution of Network Function Virtualization

Network Functions Virtualization (NFV) is likely the highest-impact technology shift in telecommunications this decade, for both CSPs and traditional vendors. Certainly, NFV — and the closely related adoption of software-defined networking (SDN) — represents the largest strategy shift in telecom since the migration from analog services to digital services.

From the early "stewardship" by the European Telecommunications Standards Institute (ETSI), commitment and buy-in from large CSPs that signaled the future will be software, to the large-scale availability and general proliferation of software-based VNFs, network operators have transformed the landscape from the central office to the datacenter.

This enormous shift has driven manufacturers and vendors to begin to move away from proprietary, sometimes-cumbersome architectures towards embracing large-scale, community-driven open source initiatives such as OpenStack.

But this shift brings with it uncertainty for all parties: monolithic, proprietary architecture-based products tend to have much tighter vertical quality assurance processes around components, design, and certifications (e.g., Apple’s successful consumer products). By opening up the network function to commercial off-the-shelf (COTS) components, the performance, service level agreements (SLAs), and reliability guarantees of closed systems, as well as the related support and service contracts, are now distributed across many third-party components.

The result is that hardware infrastructure, host environments, management and orchestration (MANO), element management systems (EMS), and hardware dependencies now must coexist across a complex and sometimes dynamic environment provided by a long list of vendors and organizations.

These new architectures pose several challenges: the multitude of combinations of hardware components, supplied by different vendors, deployed on varied software stacks (both open and closed) or on proprietary forks of these software stacks, makes striking a balance difficult. For example, outright packet processing performance and maximum compatibility are two goals that often run at odds with each other. Performance can be maximized on a targeted subset of hardware components, which limits interoperability and openness; compatibility can be maximized at the expense of optimized performance.

It was with this backdrop that I gathered together a group of world-class technology partners in a lab in Texas to see if an elastically scaling DDoS solution could be built to meet the needs of a tier-1 CSP.

Building the Ecosystem

The general parameters of the challenge were:

- To concentrate the security functions into the existing/established datacenters and to dynamically re-route/forward any inbound and outbound traffic identified for DDoS mitigation through either the closest or most available datacenter using routing-based multi-homed services that leverage the routing infrastructure for service availability
- To be built upon a fully standard NFV architecture, to run on COTS hardware, to be available on-demand, to be repurposed on-demand for multiple simultaneous functions that may co-exist (beyond DDoS security), and to be capable of self-provisioning from the infrastructure up to service through a MANO stack; these requirements go well beyond simply having compute nodes available

Network Elements	Description
VNF	DDoS detection and scrubbing solution to mitigate DDoS attacks.
Traffic Steering Function (TSF)	Intelligently steer only the DDoS traffic to the available VNFs. The Traffic Steering Function is dynamically provisioned to meet the bandwidth and packet forwarding rate requirements imposed by a sudden DDoS attack on the network.

COTS Hardware	Rack scale infrastructure, based on hyperscale principles, provides compute, storage, networking, power and cooling, and open management in a pre-integrated rack. Management is provided at the rack level and is based on the Distributed Management Task Force (DMTF) Redfish specification – industry standard open management APIs that ensure interoperability with heterogeneous systems.
MANO	Provides the auto-scaling framework that allows the VNFs to scale as required by the attack load.

Figure 1 (on the next page) shows the network topology used in this demonstration, overlayed on a simplified version of the CSP's network:

- A Security Pod is where the attack mitigation takes place and this structure is repeated throughout the network, with each of the CSP's datacenters housing a Security Pod
- This implementation relies on Anycast to route traffic to the Security Pods. With this approach, once traffic is identified as needing to go to a Security Pod, routers send the traffic to the least-cost destination pod
- Anycast also automatically accounts for datacenter availability
- Because we rely on Anycast, each Security Pod has the same network address (i.e., *lo:0)
- The different network Areas (e.g., Area0, Area1, etc.) represent OSPF (Open Shortest Path First) logical roles
- Net0, Net1, etc. are routable networks
- The large/thick, curved arrows represent Anycast tunnels that represent the book-ended flow of malicious traffic; with this book-ending, malicious and non-malicious traffic is routed into a Security Pod, and the non-malicious traffic is allowed to continue to the ultimate destination, in accordance with the challenge parameters

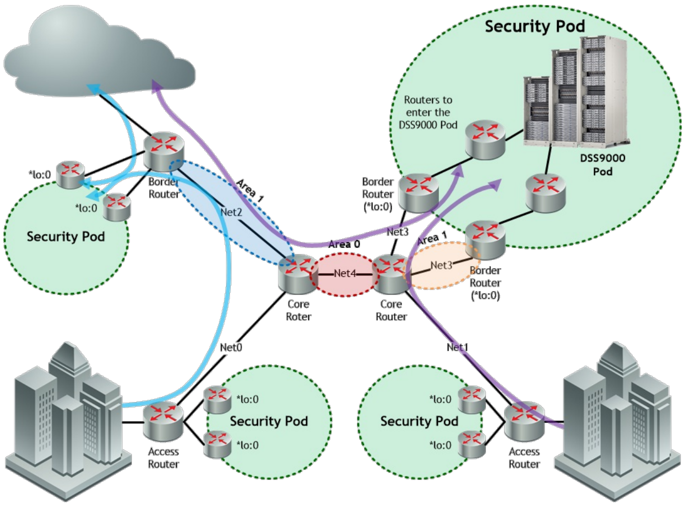

Fig. 1 - Elastically Scalable DDoS Scrubbing

The configuration for the major network elements was as follows:

Element	**Configuration**
VNF	Configured to scrub network traffic by detecting and blocking flow flood attacks:

Traffic Steering Engine (TSF)	Configured to steer all traffic to the available VNFs
MANO	Configured to scale up a new scrubber (VNF) for every 10,000 new flows per second of network traffic

Prior to the proof of concept, the network was in a steady state with a mix of traffic spread across a small number of flows. In the steady state, the number of ‘background’ flows fluctuates between about 20 and 30; the active flows represent a mix of traffic types and are of no specific importance to the DDoS scrubbing demonstration itself.

In the first part of the demonstration, a flow flood attack is triggered with a rate of 4,000 flows per second, with a flow timeout of one second; Figure 2 shows the 4,000 attack flows added to the network’s background traffic.

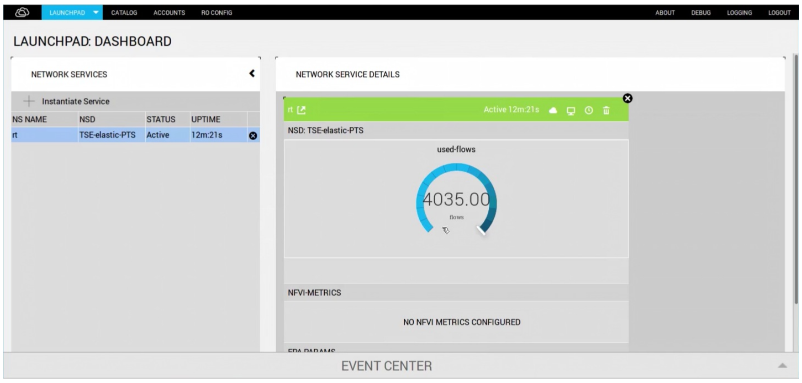

NS NAME	NSD	STATUS	UPTIME
rt	TSE-elastic-PTS	Active	12m:21s

NSD: TSE-elastic-PTS

Fig. 2 - Launchpad: Dashboard showing the small attack (4,000 flows per second) during the first part of the demonstration

A real-time look at network activity observed during five-second intervals is shown in Figure 3, demonstrating 20,000 attack flows (4,000 flows/second x 5 seconds) observed and mitigated (i.e., scrubbed).

Since this attack volume is capably handled by a single VNF instance, the MANO element has not yet triggered a second instance.

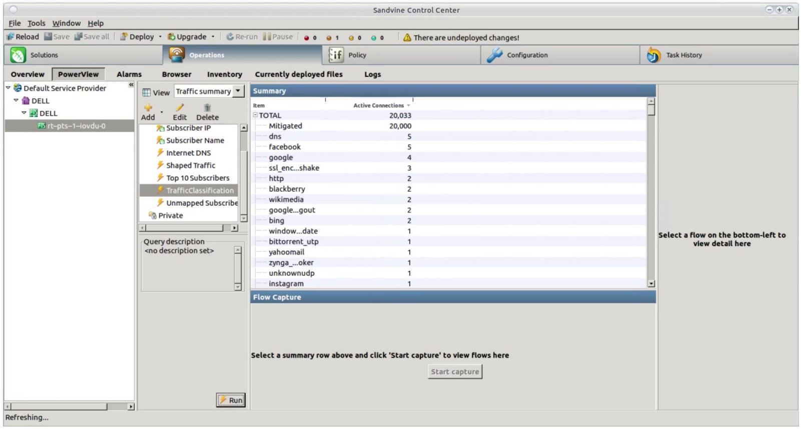

View: Traffic summary

Query description
<no description set>

Item	Active Connections
TOTAL	20,033
Mitigated	20,000
dns	5
facebook	5
google	4
ssl_enc...shake	3
http	2
blackberry	2
wikimedia	2
google...gout	2
bing	2
window...date	1
bittorrent_utp	1
yahoomail	1
zynga_...oker	1
unknownudp	1
instagram	1

Select a summary row above and click 'Start capture' to view flows here

Select a flow on the bottom-left to view detail here

Fig. 3 - Live network data show the small attack during the first part of the demonstration; this summarizes a five-second window, so the 4,000 flows per second is represented as 20,000 observed attack flows

For the next part of the demonstration, the attack scale is increased to 12,000 flows per second (Figure 4).

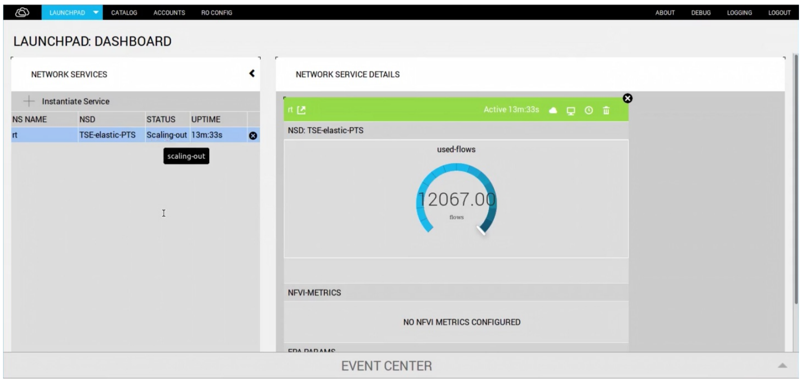

NS NAME	NSD	STATUS	UPTIME
rt	TSE-elastic-PTS	Scaling-out	13m:33s

scaling-out

rt

NSD: TSE-elastic-PTS

NO NFVI METRICS CONFIGURED

Fig. 4 - Launchpad: Dashboard showing the 12,000 flows per second used in the second stage of the demonstration

For this demonstration, the MANO element is configured to scale up a new VNF instance for every 10,000 flows; Figure 5 shows that the 12,000 flows per second attack has caused another instance to begin scaling. In the meantime, the single initial VNF instance is able to handle the larger attack.

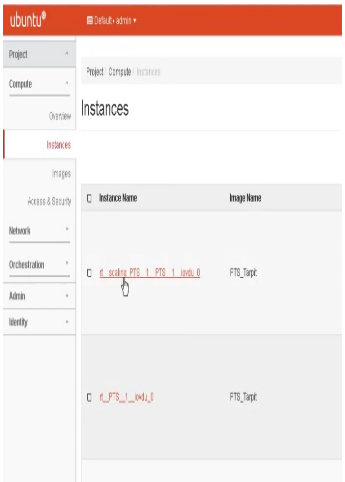

Instances

□ Instance Name	Image Name
□ rt_scaling_PTS_1_PTS_1_iovdu_0	PTS_Tarpit
□ rt_PTS_1_iovdu_0	PTS_Tarpit

Fig. 5 - Ubuntu administration panel showing another PTS instance being scaled up

Live network data show the 60,000 flows observed in each five-second interval and detection of the new VNF instance (Figure 6), as well as operation of the Traffic Steering Function. Once the VNF is online, the TSF balances the attack load between the two available VNF instances.

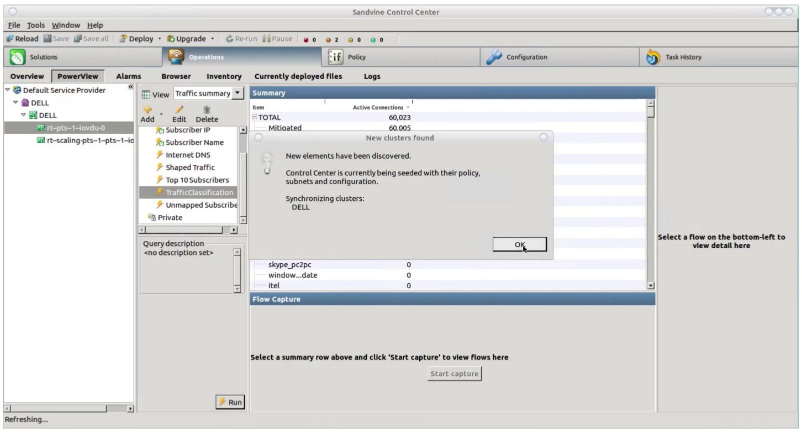

Query description
<no description set>

Item	Active Connections
TOTAL	60,023
Mitigated	60,005
skype_pc2pc	0
window...date	0
itel	0

New clusters found

New elements have been discovered.

Control Center is currently being seeded with their policy, subnets and configuration.

Synchronizing clusters:
DELL

Select a flow on the bottom-left to view detail here

Select a summary row above and click 'Start capture' to view flows here

Refreshing...

Fig. 6 – Network data show the 12,000 flows per second-attack used in the second part of the demonstration; the network has automatically detected the new VNF instance.

This iterative dynamic-scaling approach is easily extensible to DDoS attacks of any size, whether it is the “average first-half-2016” attack of 1Gbps, or the massive attack in the same period that peaked at 579Gbps.

Elastic scaling of VNFs allows for protection that is always the right size for an attack, enabling a brand new “pay-as-you-protect” model that could be much more efficient, cost-effective and highly disruptive to the current economics of network protection.