Independent market research and competitive analysis of next-generation business and technology solutions for service providers and vendors

HEAVY READING
WHITE PAPER

A New DDoS Protection Architecture for Volumetric Attacks

A Heavy Reading white paper produced for Nokia Corp.

NOKIA

AUTHOR: PATRICK DONEGAN, CONTRIBUTING ANALYST, HEAVY READING

INTRODUCTION

The information and communications technology (ICT) community is reconciled to the threat posed by distributed denial-of-service (DDoS) attacks to ICT infrastructure continuing to escalate. The cost of attacks tools available online is coming down as their sophistication is increasing. The proliferation of insecure "things" in the Internet of Things (IoT) – unstoppable in the near to medium term – promises an ever-increasing attack surface. The increasing complexity of today's networking environment means that, on their own, traditional ways of extracting visibility into network traffic are increasingly limited as tools for driving DDoS defense.

New approaches to both network visibility and network analytics, combined with greatly enhanced capabilities in network routers, create an opportunity to provide more accurate, more scalable, lower-cost DDoS protection against the stateless volumetric attacks that continue to make up the large majority of DDoS attack traffic in the network. This paper explores what that opportunity looks like from the perspective of large infrastructure providers such as telecom operators and Webscale cloud companies.

THE DDoS THREAT LANDSCAPE IS GETTING WORSE

The evidence of DDoS attacks points to the number, scale and sophistication of attacks posing an ever-greater threat to ICT infrastructure with each passing year.

- Volumetric attacks continue to account for the large majority of DDoS attack traffic directed at network infrastructure. The Verisign DDoS Trends Report for the first quarter of 2017 reports a peak attack size for the quarter of 121 Gbit/s and an average peak attack size of 14.1 Gbit/s (a 26% increase on the fourth quarter of 2016). Attacks of several hundred Gbit/s are increasingly common.
- So-called "low and slow" application layer attacks that exploit vulnerabilities in key networking protocols such as HTTP and DNS can wreak as much or more havoc, and are harder to detect, but they comprise a very small share of total DDoS attack traffic.
- Multi-vector attacks are increasingly common, such as an initial DDoS leg to preoccupy an organization's security team while a secondary attack then carries out data exfiltration. In Verisign's first quarter 2017 report, 57% of the attacks reported comprised at least two different vectors.
- And infrastructure is no longer vulnerable to just external attacks on the North-South interface or from beyond the traditional security "perimeter." Mobility, Bring Your Own Device (BYOD) in the business context, cloud, and shadow IT (employees using apps that aren't formally approved by their IT department) have punched multiple holes in the security perimeter. In addition to external sources, many DDoS and other attacks now originate from endpoints within the organization's own network infrastructure, essentially generating insider attacks embedded in East-West data center traffic.

The 2016 attacks on Dyn – enabled by leveraging the MIRAI botnet – mark a step change in the power of the DDoS attack capabilities that botnets can bring to attackers. They also represent a milestone in the sheer number of users that can be forcibly prevented from accessing online resources for hours rather than minutes.

Figure 1: Anatomy of the DDoS Attacks on DYN, October 21, 2016

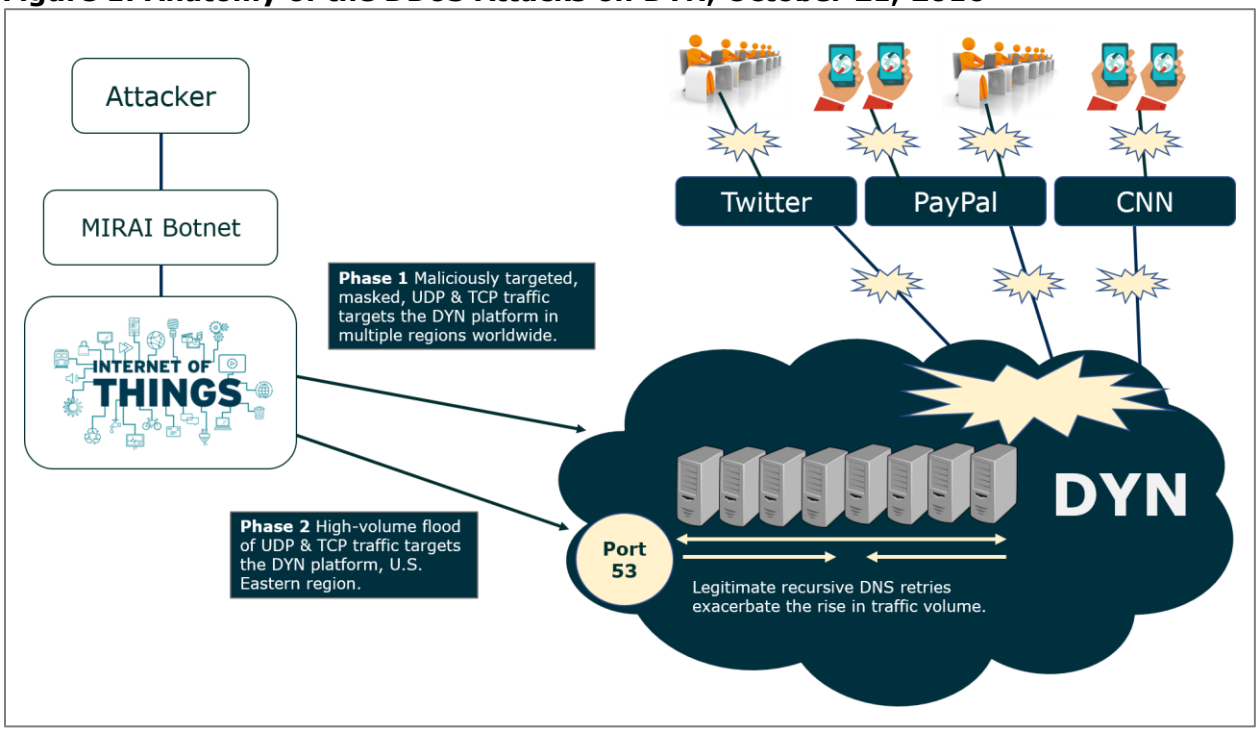

Source: Heavy Reading

Three things stand out about the capabilities of the MIRAI botnet that were used to bring down Dyn. These have become a permanent feature of the DDoS threat landscape:

- The Dyn attacks included cloud servers being hijacked and targeted to direct UDP and TCP attacks on Dyn's DNS resources.
- The endpoints exploited by MIRAI and used to generate the malicious DNS requests and create such far-reaching chaos weren't primarily PCs or servers, as used in traditional DDoS attacks; instead, they were "things" – printers, IP cameras, home gateway products – deployed as part of the IoT.
- Being created using open source software, MIRAI can be adapted and customized by attackers to create different types of attacks designed to have different impacts, as well as to evade detection.

Predictably, in October 2017 – one year following the Dyn attacks – security researchers identified a new botnet, which has been named IoT_Reaper. Unlike MIRAI which depends on leveraging default passwords, IoT_Reaper uses code exploits to enslave a number of widely deployed IoT "things." Although there hasn't been any DDoS activity from it yet, many security researchers believe IoT_Reaper is potentially a lot more powerful than MIRAI.

It Falls to Large Infrastructure Owners to Protect the Network

Enterprises typically expect to invest in their own layer of DDoS protection, including against "low and slow" application layer attacks targeting their own network infrastructure. They either buy and deploy a DDoS protection solution themselves, buy it as a managed service from third parties, or some combination of the two.

However, individual enterprises are powerless to "keep the lights on" throughout the broader Internet landscape. The responsibility for doing that inevitably falls to the large owners of infrastructure, such as telecom operators and Webscale cloud companies. These companies are unique in having the reach, the points of presence and the scale to match the scalability that attackers are able to achieve leveraging today's cloud resources and IoT footprint.

DDoS PROTECTION: THE VIEW FROM TODAY'S SOC

It's in the security operations center (SOC) that owners of large network infrastructures employ security teams to look for the first signs of DDoS traffic – either in the external environment or already within their own networks and moving laterally.

As shown in **Figure 2**, the SOC's role starts with **networking monitoring**. This consists of seeking to understand, with as much granularity as possible, exactly what is happening in the network. Most monitoring today revolves around the NetFlow data gleaned from routers.

Figure 2: The DDoS Defense Lifecycle Today

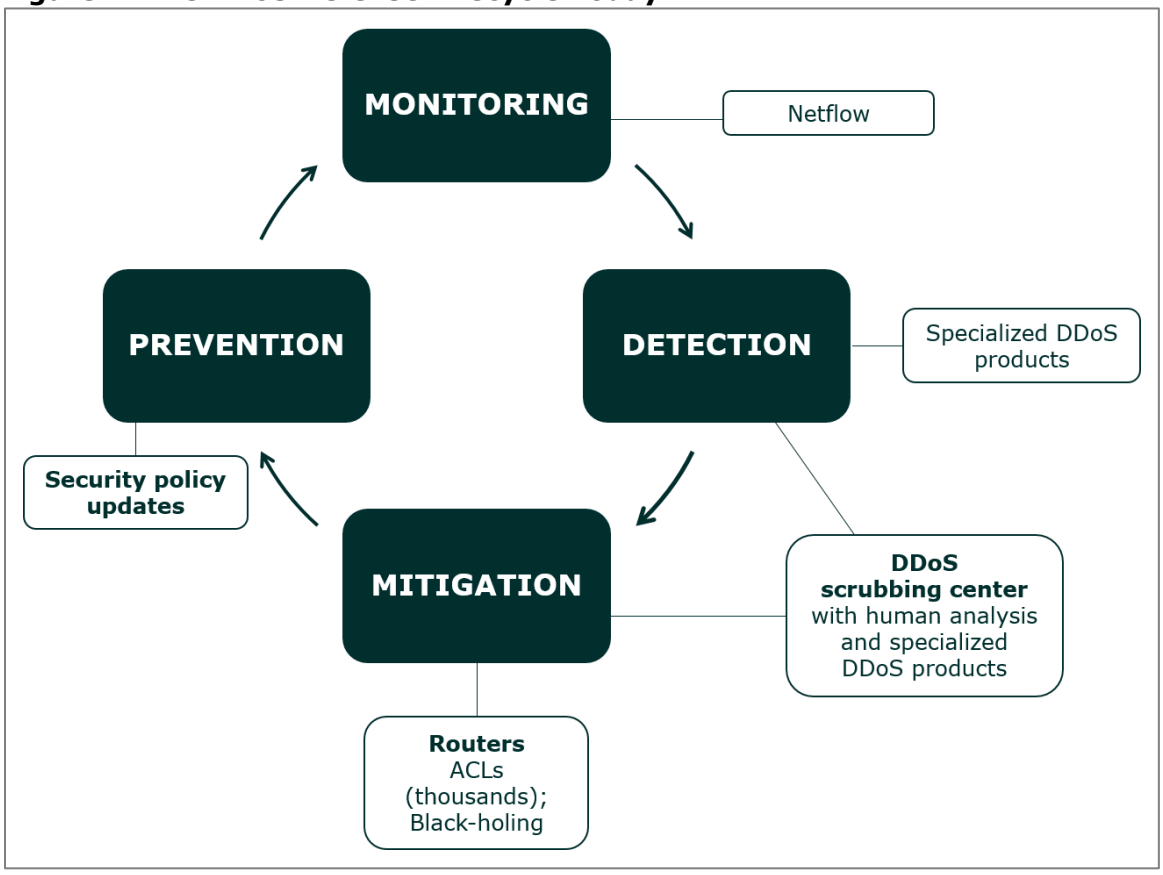

Source: Heavy Reading

The bane of the SOC operative's life today is the number of alerts generated by network monitoring that routinely prove to be wrong. False negatives – allowing malicious traffic through – are particularly harmful. But false positives – flagging perfectly good traffic as potentially bad – are more common and can be a major distraction.

When it comes to the initial **detection phase**, the faster the SOC can accurately detect malicious traffic, the better it can minimize the impact of an attack. The accuracy of detection is determined by the quality of packet inspection, the quality of contextual data and the quality of the analytics applied to the traffic.

DDoS monitoring products initially spot bandwidth anomalies and flag them for the attention of the DDoS mitigation or scrubbing center. Here, other specialized DDoS products determine whether or not the traffic is indeed malicious, using deep packet inspection (DPI). If it's bad, the SOC can choose from a variety of options with respect to the **mitigation phase** by means of blackholing the traffic, applying access control lists (ACLs) or network filters.

With the scale, reach, frequency and sophistication of DDoS attacks set to increase inexorably, operators of large network infrastructures risk facing a steep rise in their spending on defensive infrastructure in order to maintain network availability. They are looking for ways to augment the monitoring, detection and mitigation capabilities of the DDoS defense lifecycle for more accurate, faster, protection at lower cost. The next sections look at some of the new options that are becoming available.

MORE GRANULAR VISIBILITY REDUCES FALSE POSITIVES

One of the biggest challenges in detecting volumetric attacks is to determine the causes of bandwidth anomalies when they are first identified. Surges in total traffic volume – or in a particular type of traffic – that arise from entirely legitimate changes in behavior by end users or by other network infrastructure owners can look very similar to a malicious attack.

Figure 3: Increased Network Complexity Reduces Visibility

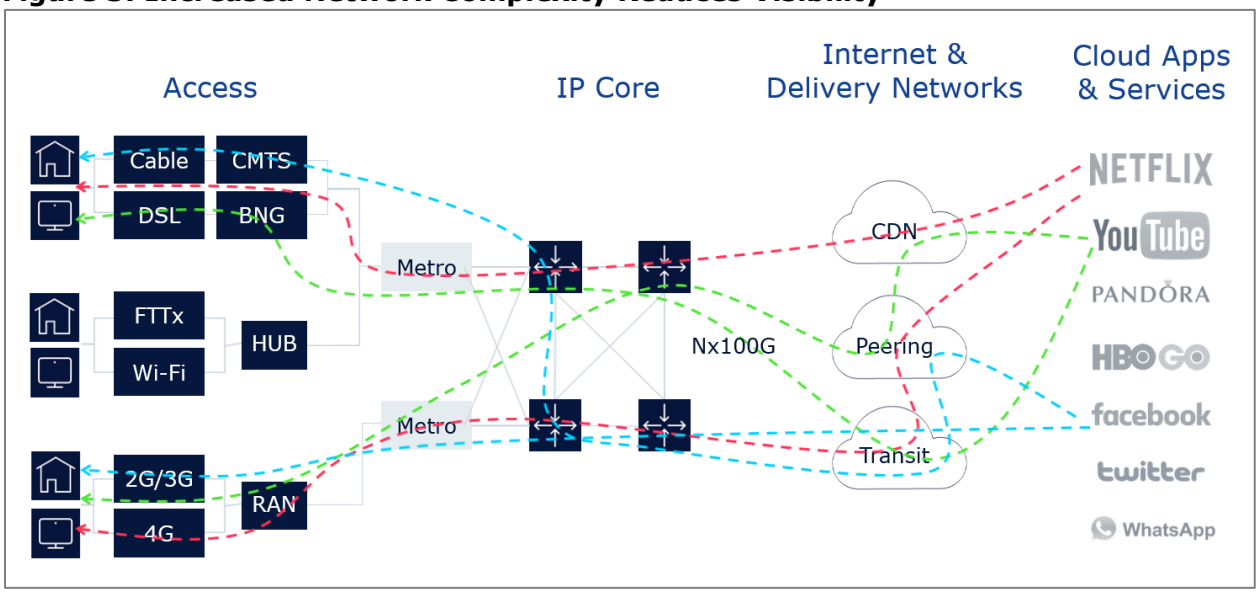

Source: Nokia

The SOC needs as much visibility into the broader network context in order to make the right call each time it is faced with a significant traffic anomaly. Some factors have actually served to make matters worse in recent years:

- Many of the largest operators report total traffic growing 40-50% a year.
- The underlying complexity of the networking environment has accelerated rapidly. Today a huge amount of traffic is delivered as streaming data delivered from the cloud, via Webscale cloud operators. That traffic also transits large intermediary actors in the networking value chain, like big content distribution network (CDNs) such as Akamai. Knowing that a packet originates from a Level 3 data center, for example, has some value. But in a DDoS protection context, not being able to see that that packet originates from a third-party video services cache from within that Level 3 environment masks what's really going on in the network. Knowing that it does enhances visibility significantly.

Three important things are changing in the networking environment that are bringing more visibility into what's going on in the network – and why.

- Huge step changes in storage and compute technologies now allow the entire supply chain of the Internet to be captured and presented in real time to operators of large network infrastructures. There is significant added value from a DDoS protection standpoint in being able to determine that a packet isn't just any old packet, but that it's a Facebook or a Netflix packet, that it has traversed this or that specific infrastructure, and that it's headed for this or that specific home gateway.
- Today's big data technology also allows colossal volumes of additional network data to be gathered. Other network data, such as DNS data and other server logs as well as threat intelligence feeds, can be brought to bear, in addition to NetFlow, in order to bring the maximum possible relevant data to bear.
- Advances in analytics algorithms make for more accurate determinations. The combination of big data and advanced analytics sharpens the SOC's understanding of traffic patterns and allows for a reduction in the number of false positives.

Visibility into the supply chain of the Internet, combined with big data and the latest analytics software, should serve to augment the detection phase in the DDoS protection lifecycle. They can be combined to ingest the vast amounts of contextual data that can now be gleaned from network monitoring – and then correlate it to enhance the accuracy of DDoS threat detection.

This should also reduce the time it takes to make accurate determinations. That's important because in the time spent between identifying traffic as potentially suspicious and then formally confirming it as such, further damage is liable to be inflicted.

Such analytics capabilities need to run on industry-standard servers to keep costs down. They also need to support open application programming interfaces (APIs) so that they can easily integrate and apply analytics algorithms to vast amounts of network-related data and threat intelligence feeds generated both by the infrastructure operator itself as well as brought in from specialist third parties.

The addition of analytics dimensions is also important now in the context of the increasing share of network traffic that is being encrypted. Decrypting, inspecting, and then re-encrypting traffic for the purposes of threat detection is expensive, and is typically done by DPI products in hardware. By contrast, a great deal of contextual analytics occurs independently of any encryption.

PACKET INSPECTION IS INCREASINGLY ABOUT "WHERE?"

Unlike in the big data and analytics space, there hasn't been a great leap forward as such in the field of IP packet inspection. From a DDoS protection perspective, discussion around packet inspection tends to revolve around the blending of packet inspection and analytics outcomes in pursuit of faster time to accurate detection of DDoS threats and, increasingly, on where packet inspection should be carried out.

Where in the network the detection is carried out, and on what platforms, is important in terms of both cost and security efficacy. As depicted in **Figure 2**, the conventional DDoS protection architecture typically sees anomalous traffic systematically forwarded to a centralized DDoS scrubbing or mitigation center, where specialized DDoS protection products perform packet inspection on the suspicious traffic.

There are several issues with this conventional approach:

- Forwarding traffic to a centralized location for treatment delays the treatment compared with protecting at the edge when the traffic first enters the network. Delay carries the risk of additional harm.
- In the case of the large volumetric attacks that compromise the large majority of DDoS attack traffic, forwarding suspicious traffic to a scrubbing center also carries a significant cost in terms of network bandwidth.
- Scrubbing centers use DDoS protection products that are purpose-designed for the DDoS protection applications. As attacks begin to cross the n x 100 Gbit/s and even terabit boundaries, standard network elements designed for high-speed processing, such as switches and routers, have potential to offer a more cost-effective approach to responding at scale.
- The direct cost can be very high in the common case, where the mitigation service is charged for according to the volume of data traffic. That means the cost incurred is very high in the increasingly common case of large-scale bandwidth attacks at or approaching 1 Tbit/s.

The conventional approach tends not to distinguish between application layer attacks and volumetric attacks as clearly as is needed. While application layer attacks require stateful analysis in a scrubbing center and significant human intelligence inputs, the vast majority of volumetric attacks can be detected and mitigated with automated processes that do not require stateful analysis, hence costly diversion to a centralized scrubbing center.

COST-EFFECTIVE USE OF ROUTER INFRASTRUCTURE

Historically routers have tended to feature at the very start and very end of the DDoS defense lifecycle. They yield NetFlow data, anomalies in which serves to trigger the very beginning of the detection phase. They are then used to implement ACLs – typically in the several thousands – and black-holing in the mitigation phase.

Routers have tended to be excluded altogether from the core detection phase. Until now, the dearth of contextual information enhanced by analytics; moderate growth in the number of Internet end points; as well as the trade-off between capacity and performance in router

products themselves has tended to limit routers to this minor supporting role in the DDoS defense lifecycle. There's a good case for this approach to be reviewed nowadays. In particular, there's a case for routers to be used for detection and mitigation of volumetric attacks.

Figure 4: New Options in the DDoS Defense Lifecycle

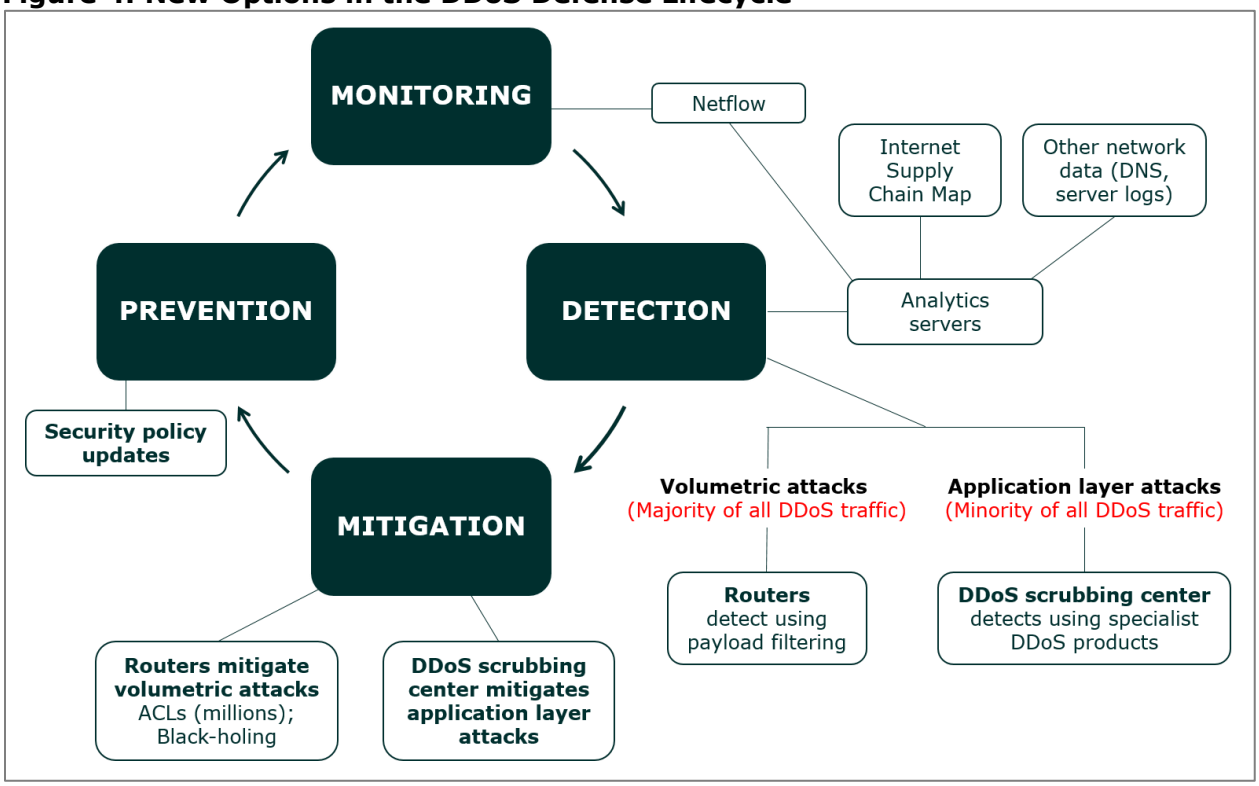

Source: Heavy Reading

The role of analytics and the huge growth in vulnerable IoT end points as potential game-changers has already been discussed. In addition, routing technology has undergone one of its periodic step changes in capacity of late.

In some cases, new network processors should now be able to allow routers to perform payload filtering as well as run literally millions of DDoS attack filters without materially impacting their performance. And routing infrastructure isn't just more cost-effective than specialized DDoS detection; it's also more scalable than most of today's DPI-based DDoS product platforms, some of which can be more hardware-dependent for scalability.

Don't Throw Out the Baby With the Bathwater

There's an important rider to the above: Real-time correlation of multi-dimensional data combined with modern routing infrastructure is well suited to identifying and mitigating most volumetric attacks, which comprise the large majority of DDoS attacks.

Application layer attacks, on the other hand, tend to be more difficult to detect and require stateful analysis to do so. This subset of total DDoS attack traffic remains best suited to the conventional DDoS protection model – direction toward a mitigation or scrubbing center for the operator's security team to undertake further analysis and investigation leveraging traditional solutions.

SUMMARY

DDoS scrubbing centers have long featured at the heart of DDoS protection strategies for owners of large network infrastructures, such as telecom operators and Webscale cloud companies. They continue to play a critical role in detecting and mitigating sophisticated application layer attacks that are stateful.

For the large majority of DDoS attacks that are volumetric and stateless, owners of large infrastructures need to evaluate the combination of new capabilities in big data analytics and router platforms that have potential to improve the efficacy and reduce the costs of DDoS protection.

ABOUT NOKIA

Nokia (www.nokia.com) serves communications service providers, governments, large enterprises and consumers with the industry's most complete, end-to-end portfolio of products, services and licensing. From the enabling infrastructure for 5G and the Internet of Things, to emerging applications in virtual reality and digital health, Nokia is shaping the future of technology to transform the human experience. A truly global company, Nokia's 160 nationalities work in more than 100 countries.

Nokia Deepfield Core Platform is a big data engine/software platform that enables service providers (cable providers, cloud providers and telcos) as well as large enterprises and Webscale companies to understand their IP network with unprecedented visibility, all in one place. Nokia Deepfield Cloud Genome is a patented map of the global service supply chain that adds visibility to all applications built onto the Core Platform.

The big data engine/software platform is the basis for Deepfield analytics applications, including Cloud intelligence that provides end-to-end network visibility, Service Intelligence that monitors customer QoE in real time, and Deepfield Defender, a security application that performs real-time distributed denial-of-service (DDoS) detection and mitigation.