Research Note

August 2021

A better architecture to fight DDoS

To Meet Rising Attack Volumes, Efficacy and Economics must go hand-in hand

Author: Grant Lenahan, Principal Analyst

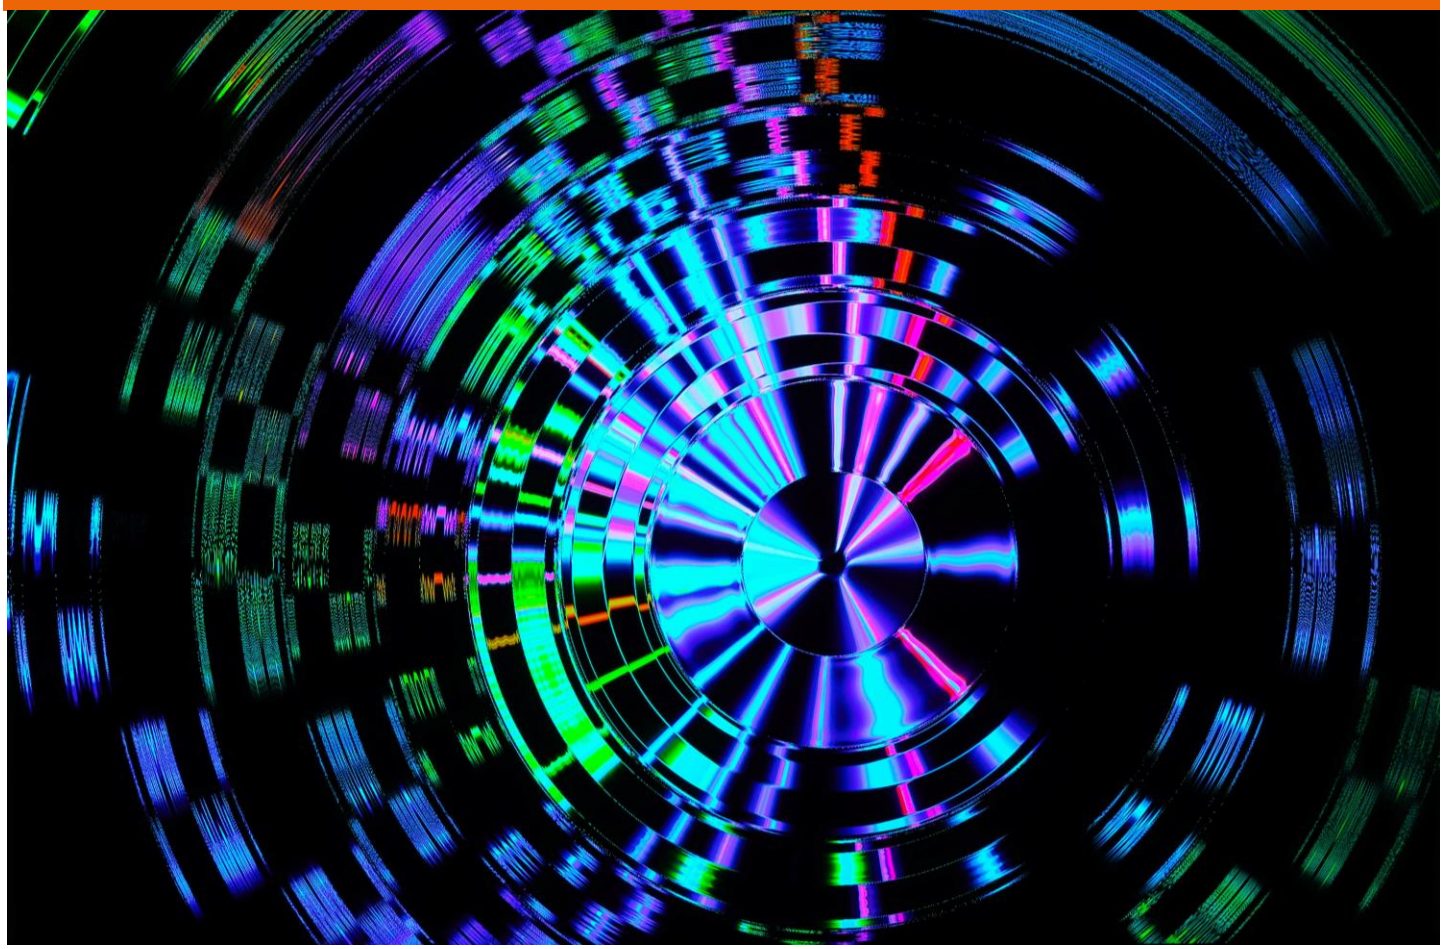

www.appledoreresearch.com

Appledore RESEARCH

Published by Appledore Research LLC • 44 Summer Street Dover, NH. 03820

Tel: +1 603 969 2125 • Email: info@appledorerg.com• www.appledoreresearch.com

Figures and projections contained in this report are based on publicly available information only and are produced by the Research Division of Appledore Research LLC independently of any client-specific work within Appledore Research LLC. The opinions expressed are those of the stated authors only.

Publish date: 8/1/2021

Photo by FLY:D on Unsplash

OVERVIEW AND SUMMARY

Data indicates that success demands a new approach to DDoS

DDoS and other techniques to execute ransomware are on the rise, costing everyone from individuals to huge enterprises significant cost, inconvenience and lost revenues. For some smaller online enterprises, there is ample evidence that DDoS may cost them their business. Appledore believe that cyber security in general and DDoS/ransomware protection in particular are large and realistic opportunity (as opposed to some of the industry's past "pipe dreams") for network provides to grow and diversify their revenue base.

Recently, Nokia Deepfield shared the results of an extensive analysis of DDoS attack data collected by their own Deepfield solution across operators (with all data anonymized, of course). This data is fascinating in that it shines a light on the sources of DDoS attacks and repudiates many common themes circling today about new, different, sophisticated methods. According to the data, DDoS is not changing, nor are its vectors many and varied. It's simply scaling up, recruiting more IoT devices as foot soldiers, and becoming a more dangerous and costly business, using the same methods and effectively exploiting an economic/volume mismatch between attackers and the attacked.

Concurrent with this study, Nokia is promoting a new[1] approach to DDoS protection and remediation. In the "narrow" Nokia demonstrate the efficacy and cost efficiency of an in-line solution based on Deepfield analytics, FP4-based intelligent Service Routers, and both real-time (execution) and slow (analytics/learning) control loops. In the large, they blueprint a multi-vendor approach that has significant cost savings and protection benefits, along with a parametric economic analysis of traditional (PMO) and future (FMO) approaches.

Appledore commented on the promise of Nokia's approach, which maps to Appledore's generalized next-generation taxonomy, as early as the FP4 launch in the summer of 2017. We continue to believe that solutions structured with ML to perform the "slow" learning loop, coupled with rules and masks to effect the "fast" or corrective loop, are the correct approach for myriad tasks. This is one. The right design is important, but the real value lies in the cost profile; which exhibits vastly better scale economics. This is critical if DDoS volumes continue to rise, and if SPs wish to aggressively market network-based protection services, without undue concern for the incremental cost, and confident in the service' margins.

In this note, we review and comment on the concept, the fundamental economic promise, the Nokia implementation, and the implementation beyond Nokia's. We believe that highly automated solutions with good scaling performance and minimal human intervention are the best path to both security and cost effectiveness/scale.

[1] Actually, it is an existing approach, but relatively unused until very recently.

Finally, we believe that, in order to stop attacks from overwhelming networks themselves, the only logical delivery entity for this service is the Service Provider Community – and therefore it is a significant business opportunity for the industry.

KNOW THE ENEMY: DDOS CHARACTERISTICS, TRENDS AND TRAJECTORY

DDoS attacks have been exponentially increasing in number and damage over the years, driven mostly by monetization. Basically, attackers ransom the attacked, or they have an interest in talking the attacked offline for political or business reasons. The mainstream story line makes it appear that there is significant technical evolution, growing sophistication and diverse sources of these attacks, paving the way for more sophisticated solutions. Nokia's Deepfield data tell a meaningfully different story however – and, if understood, presents a good understanding around which to build effective and cost-effective solutions using existing technologies.

Over the past months, Nokia Deepfield has collected a large corpus of data on DDoS sources, methods and attacks across many service providers (data anonymous of course). According to Deepfield's data, DDoS is in fact growing rapidly – in fact exponentially. The cost is similarly growing rapidly. But the methods have NOT changed dramatically, and the sources of attack origin are surprisingly limited – as few as 100 domains account for the vast majority of DDoS attack originations and can be identified in advance. In fact, Deepfield claims to have compiled that list and to be keeping it up to date, via its proprietary Secure Genome approach. A public NANOG presentation (video) by Dr. Craig Labovitz with more information on attacks and methods may be found here.

DDoS methods are relatively stable as well and limited in number. All rely on sending small volume queries to public servers that amplify this traffic with responses that are known to be many-fold larger (as much as by 10,000X!). By exploiting such asymmetry, the economics work in the favor of the attacker. The major methods are 1) IP address spoofing to overwhelm firewall state or generate large amplified bandwidth floods or 2) botnet attacks against applications. The simplified goal remains to initiate a small, cheap query and send the lopsidedly huge response to the target address. Clearly one response must be the ability to level economics and handle large, growing and regular volumes more cheaply than is done by today's DDoS methods.

The motivation for DDoS is also generally consistent: monetary gain. The manifestations are more surprising. We would expect extortion to be the leading player, and it is clearly top 3. But the other two are gaming and gambling – where the objective is to disadvantage or even cripple your opponent and therefore win money. Can high speed market trading, which depends on millisecond advantages, be far behind? Its largely gambling after all.

TODAY'S MOST COMMON APPROACH AND ITS LIMITATIONS:

Setting the stage, it is worth remembering that flows over the internet typically transit several different service providers from their origination to destination and pass through one or more peering points. Today's DDoS solutions are primarily reactive. First, detection is left up to the end service provider or the attacked entity to identify and begin corrective actions. Second, these attacks are typically detected only after they begin to scale – with scale being one of the identifying characteristics. Yikes.

The majority of network-based DDoS remediation solutions have historically been – and even today are – based on centralized scrubbing centers where dedicated scrubbing appliances are capable of removing DDoS traffic and returning valid traffic. A simplified diagram of such an approach is below. These centralized resources demand, by definition, dedicated back-haul network capacity from the point of detection to the location of the closest "scrubbing center", and, for the clean traffic, back (important). Finally, since routers up to the scrubbing center do not have the ability to distinguish valid from DDoS traffic in a mixed stream, nor to handle it independently, they must route and back haul all of the questionable traffic to a scrubbing center. This typically means all traffic destined for the monitored IP prefix(s) from a particular ingress point. Note that if distributed IoT devices are employed for amplification – they may originate in many areas including as local devices to the network provider.

Figure 1: Simplified functional architecture / path of traffic handled by scrubbing centers

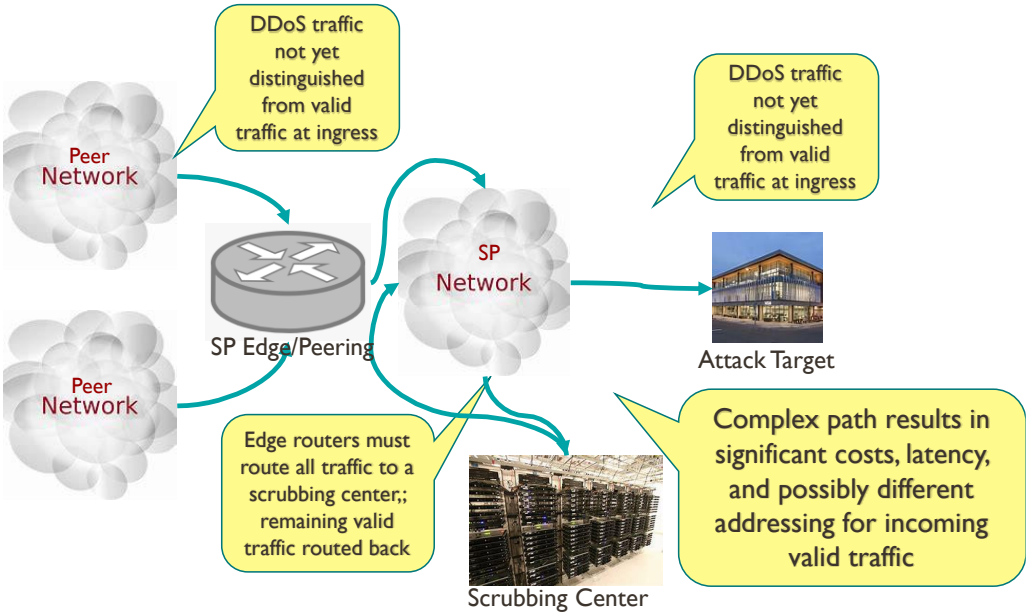

Source: Appledore Research

Such scrubbing center-based solutions exhibit poor scale economics; demanding *more than* 1:1 dedicated back-haul capacity plus DPI/scrubbing capacity for every megabyte of attack. Note that this means "capacity over and above that required for normal traffic handling". By creating very

costly situations, DDoS attackers can win even if they ultimately "lose". This also imposes a practical cap on DDoS capacity at any given moment – and any DDoS above that installed capacity cannot be remediated. It is critical that service providers not only have a solution that is cost effective at pre-existing DDoS traffic volumes, but one that can scale. Otherwise, attackers can, in effect, lose the tech war and win the economic war – and worse, they know it, eliminating any deterrent value.

To a large degree this situation is a function of router economics and capacity. Routers and their core chips are designs to route, not DPI/scrub. Just as other applications of DPI (QoS, metered charging) are highly targeted, so is DDoS scrubbing capacity. And since Scrubbing equipment is not the same as in-line routing equipment, and since no one knows where the next attack will come from, it tends to be centralized in a few places, and back-hauled to those centers. This further worsens the economics of scale since not only must routing/scrubbing capacity grow proportionally with attack volume and overall traffic (typically 15%-20% of overall traffic and growing rapidly), but transport must now be back-hauled from ingress to a shared "scrubbing" center where dedicated hardware performs packet inspection, removes DDoS traffic, and re-sends legitimate traffic. The result is two costs that scale linearly with attack volume: 1) scrubbing hardware and 2) back-haul capacity.

IMPLICATIONS FOR DDOS SOLUTIONS

DDoS methods are not new and different, they are just more, larger and better (more below). The observed characteristics and direction in DDoS attacks suggest that what is needed is not a widely diverse, new set of protections, but rather solutions that have five characteristics:

1. Scales affordably in the face of large and growing data volumes
2. Identifies sources and attacks early
3. Identifies and remediates attacks as close to the source as possible
4. Blocks traffic before it paralyzes targets – or the network segments proximate to them
5. The ability to translate spoofed and other "interim" addresses back to known malicious sources
6. Automation, to both speed remediation and deal with scale in the face of a growing industry expertise shortage

All of this implies a solution that proactively identifies likely sources and fingerprints of attacks, and sees through spoofing and other methods of obfuscation, such that they can be blocked early and locally. It also suggests an architecture that does not demand building huge capacity to meet the need; performing costly back-haul, and therefore engaging in a lopsided war of attrition with an enemy that may have a 100:1 or 1000:1 (or more) advantage based on the amplification inherent in the methods employed.

According to the data presented, remediation is less technically complex than typically presumed, and simpler than the common narrative implies. Attacks tend to originate in a relatively small number of sources, which, amazingly, advertised their services – sometimes whitewashed as "for testing against DDoS". The chart below (courtesy: Nokia) shows the chain from origination, through

amplification and amplified response, resulting in a multi-terabyte flood intended to simply flood and overwhelm the target. This implies that, if we can identify a relatively small number of origination points, and trace spoofed addresses as they move through the network, DDoS can be blocked before it causes significant damage. This blocking, of course, must be affordable and the 1000+:1 advantage held by attackers means solutions must be carefully engineered for affordable scale.

Figure 2: DDoS methods and originations are limited - and can be identified

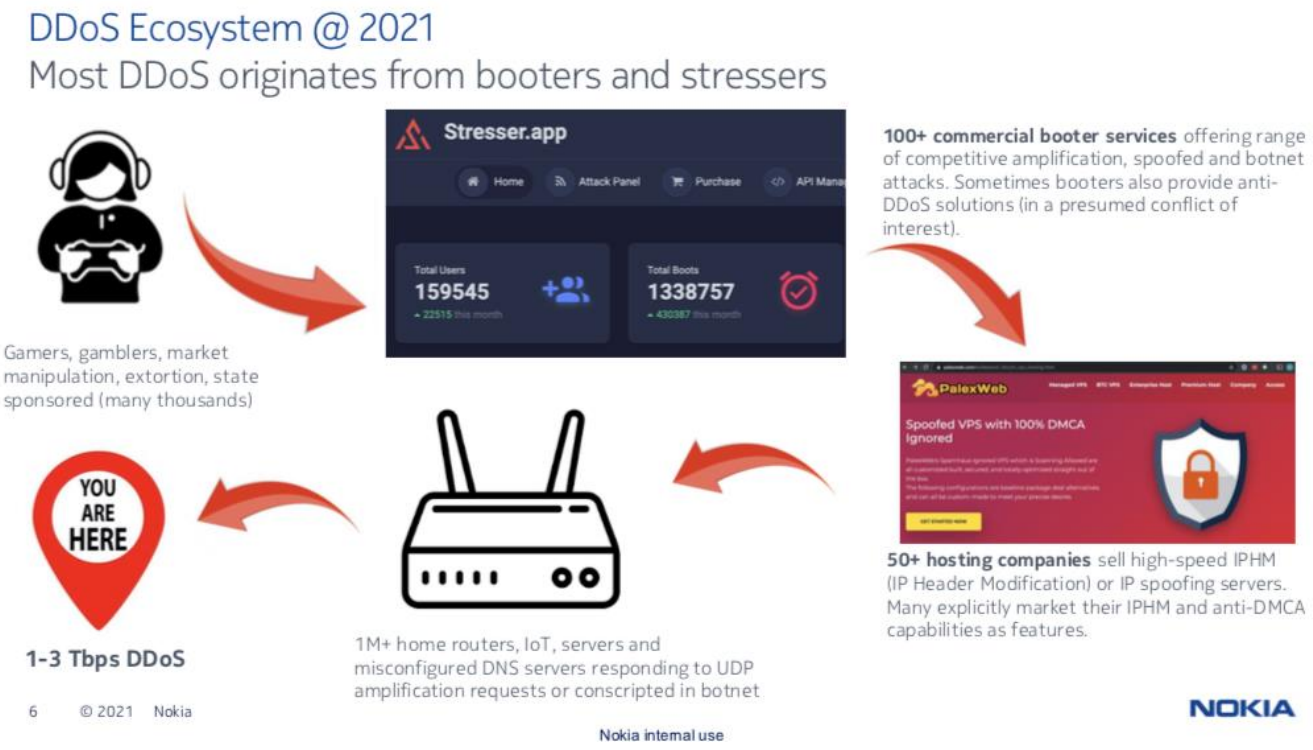

Courtesy: Nokia Deepfield

HOW TO WIN: ENGINEER BETTER ECONOMICS AT SCALE

Appledore have advocated for intelligent, in-line router architectures for some time. They represent a slightly higher initial cost (for the ability to do traffic inspection, e.g..: processing capacity at high scale using Flowspec, Netconf, other . . .) and to do so without impacting the traffic throughput rate. Installing such capacity, therefore represents an investment in the future. The tradeoff will be higher cost today, against lower costs of handling DDoS tomorrow. While every implementation and set of vendors will have a slightly different cost and break-even point, we believe that the general form of the comparison is relatively consistent. Data (again from Nokia) correlated by several operators, using both Nokia and other hardware solutions, confirms that the incremental cost of handling a

DDoS attack may be reduced by 75% to 90% depending on the traffic assumptions employed. The form of the cost savings is depicted (illustratively) below.

Figure 3: Relative Costs (see notes below)

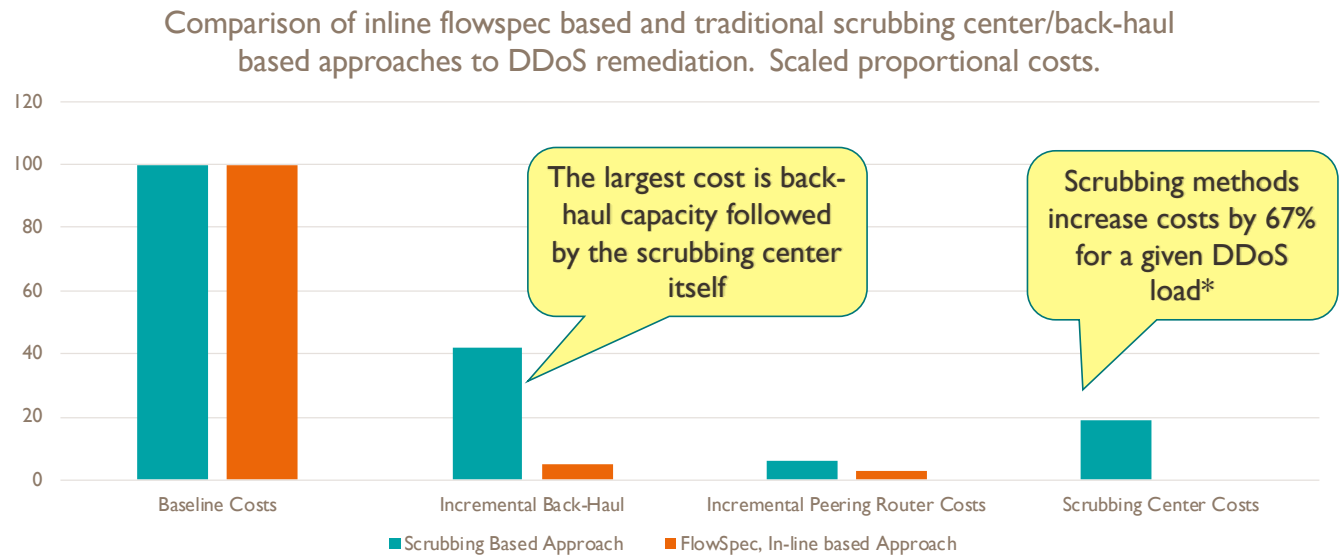

Source: Appledore Research (Data: courtesy Nokia Bell Labs)

** Costs are parametric with baseline traffic, scrubbing center locations relative to attacks, and peak DDoS incremental traffic. This example uses a 5Tb/s baseline traffic load, with a 2 Tb/s DDoS spike. Reader should assume that current growth trajectories will continue, making the economics presented even more compelling, and the threat even more challenging*

BEYOND LOWER COST – A RECIPE FOR A BETTER APPROACH

An improved DDoS solution architecture needs three characteristics:

1. Elimination or vast reduction in back-haul
2. Earlier and more sensitive identification of threats and threat sources
3. Ability to cleanse in-line traffic at the rate of the total offered load; and the ability to scale cost effectively with attacks.

These requirements fit a simplified functional approach that Appledore has advocated for ages. In its generalized form it is the basis of our Appledore Network Automation Software taxonomy – a reference that we apply to myriad use cases. Further data can be found here.

Figure 4: Appledore Management Taxonomy Focuses on Automation and Structured, Closed Loops

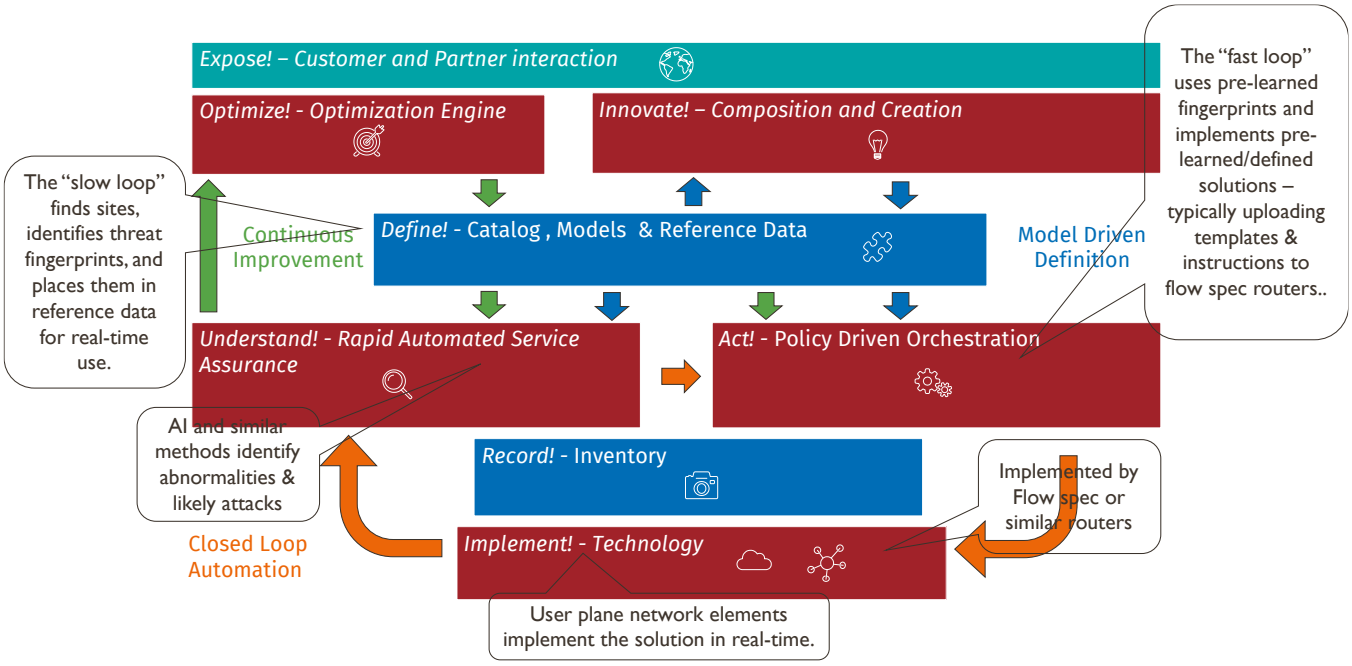

Source: Appledore Research

A few points are worth mentioning. First, Appledore's taxonomy forms a giant, multi-functional closed loop in which data is collected, normalized, analyzed, and then actions are calculated and implemented. As a closed loop, it has the benefit of being self-correcting and if desired, learning. Second, there is often both a "fast" and a "slow" loop. AI and large data sets, almost by definition, are not real time (how does one trend over time, in real time?). Yet these "slow" loops perform the magic to identify fingerprints and track the efficacy of various solutions – constantly placing the best and most up-to-date information in reference stores for real-time use.

Figure 5: Fast and Slow Loops Applied to DDoS Identification and Remediation

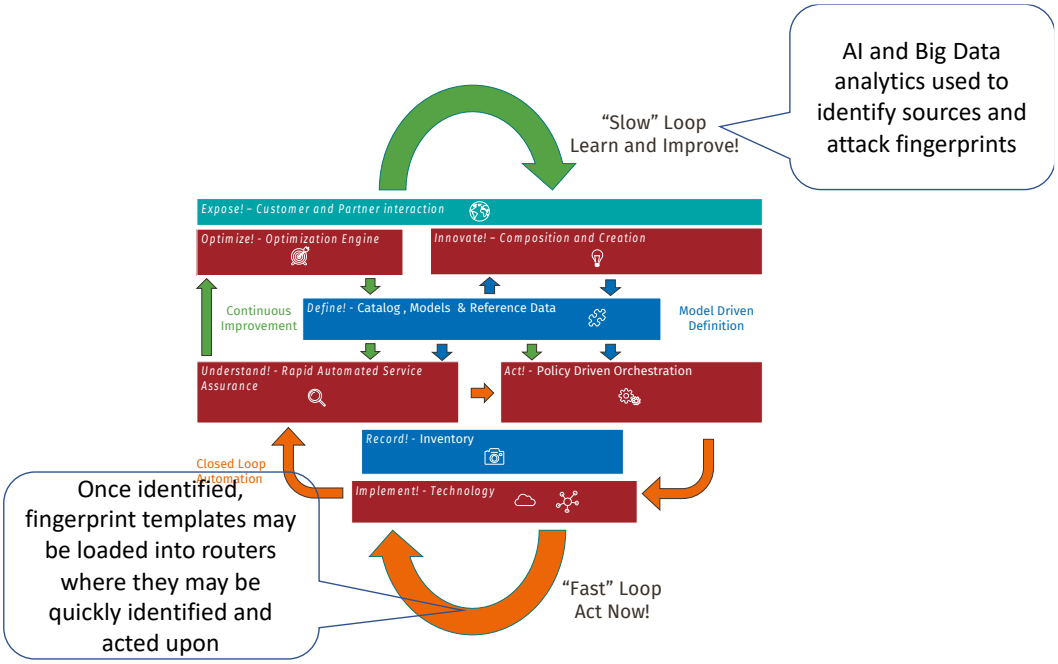

Source: Appledore Research

The Nokia proposed approach, and the associated Bell Labs' cost analysis, map to this taxonomy. No doubt individual vendors' approaches will have small variations on the theme – maybe a learning process in action that will result in continued refinements.

DDOS IS A NETWORK PROVIDER BUSINESS OPPORTUNITY

As with most cancers, early detection is one key to successful remediation. Given the fact that most DDoS originates from a few sources, and that its amplification points can change dynamically and come from many ingress points, only the network provider can identify and potentially arrest an attack before it begins to clog the network (at best) or completely overwhelm segments (worse) and finally bring down the applications/users under attack (the ultimate goal). We see several reasons why network providers are the logical – maybe necessary – point of detection and remediation:

1. Only network providers own and operate the edge routers capable of inline suppression
2. Only network providers can continuously observe abnormal behavior across ALL ingress points and from "things" connected across the network.
    a. *Appledore comment: broad based inspection is required to counter the "Distributed" aspect of the DoS attack. Inspecting on one spot is not enough*
3. There are large economies of scale and scope to a network provider/operator continually collecting data on normal behavior, and therefore being able to characterize abnormal behavior
4. There are also large economies of scale in shared remediation hardware

For these and other reasons, Appledore believe that DDoS remediation, especially when in-line edge Netconf and Flowspec capable routers are employed, is a natural activity and therefore business for network providers. We already see several major CSPs offering versions of such as network-provided DDoS protection and believe that, if packaged and priced accordingly, it is one of those areas where the basic economics and efficacy profiles allow simultaneously for solid margins (to the network provider) and significant TCO savings for the enterprise/targets.

The industry is already picking up on this opportunity and approach. Deepfield claims that more than 10 customers, including StarHub and Exponential-e, have selected Deepfield Defender for securing their networks and infrastructure from DDoS threats and delivering network-based DDoS solutions to their customers. Appledore has also discussed with others CSPs similar solutions, some of which include alternative vendors (one public case is Lumen, employing Arbor).

CAN A COLLECTIVE INDUSTRY EFFORT NIP ATTACKS AT THE BUD?

We have discussed the advantages of identifying an attack early, and terminating it immediately after it enters a network provider's domain. We have also discussed the economics of scale and scope that favor operation of DDoS protection "aaS" by network operators, for a large group of connected enterprises. Finally, we have noted that the data suggests that the number of originating locations for DDoS attacks is small and they may be identified in advance, facilitating cheaper, earlier remediation. This logic can be extended from each individual network operator to an industry collective action, and Nokia Deepfield advocates such an extension.

Industry collaboration could allow attacks to be detected and terminated wherever they originate. For example, the attack could begin *upstream*, in country A, on network X, cross a peering location, and attack an enterprise in Country B on network Y. Complicating this, "things" could be recruited on several surrounding networks – so attack traffic might come from more than one source. Presumably the attack source from country A is employed in many attacks on targets in many nations. Presumably also, those "things" used as traffic amplifiers will be recruited to attack multiple targets over time. Hold that thought for a solution that could be much better than "every network for themselves".

The best solution would be for each network provided to patrol its own borders and terminate attacks (either the origination messages or the amplified flood, or both) as early as possible – preventing it not only from reaching its target on a peer network, but preventing it from clogging the transited network(s). Such a collective industry effort would lower total costs and prevent more DDoS traffic from crossing interim networks and peering locations – all of which cost money and reduce perceived network quality via avoidable congestion.

This is a classic case of economic "externalities" as economists call them. There is a clear economic advantage in aggregate, yet the costs may sometimes be borne by one subset of networks, while the

benefits accrue to others. Therefore, some method of compensation of collective investment may be in order.

SUMMARY AND RECOMMENDATIONS

Appledore believes that the data presented by Deepfield's analysis, along with the growth in DDoS volumes and the significant economic loss to end users should send a clear signal that new, better methods ought to be employed. At the same time, we see a silver lining for the network provider industry, in that there is also a clear business opportunity to create a profitable service, in which networks enjoy a clear competitive advantage.

In related research, we have argued that the changing network environment, the escalating cyber security threat (and its recognition across enterprise and public entities) creates bot a perfect storm for better security measure and vastly different approaches – mostly based on replacing error-prone and unscalable manual methods with DevOps-led, error-free and scalable automation.

We believe that network operators ought to invest in remediation approaches that may cost a little more now, but save significantly over time. This means reducing or eliminating the back-haul and high-cost redundant scrubbing capacity and replacing it with edge routers capable of high performance, and filtering and selective routing at line rate. By making this (relatively smaller) investment, network operators can cleanse streams early, and at a far smaller incremental cost. As always, this demands investing now, and garnering the rewards later – which has long been a business case challenge in our conservative industry.

In fact, it may be advisable to think beyond DDoS alone on such investments. In-line, Flowspec/Netconf routers will have myriad beneficial uses in the network beyond DDoS, and represent a flexible, cost effective approach to intelligent traffic handling. Note: it is important to consider the throughput capabilities of those routers under a pattern-matching load – and in fact considering various loads and volumes of templates to match. It is highly likely that will vary greatly in these regards.

Remediating a problem demands first identifying and isolating that problem. Big data, proactive approaches allow easier identification of attacks – and this can extend beyond DDoS. By constantly looking at patterns, and by proactively "crawling" sites to better understand them, such approaches can speed remediation, lower costs, potentially avoid paralyzing traffic volumes, and the approach may be applied well beyond DDoS alone

We also see merit in an industry collective approach with each network patrolling its own borders. This has local benefit (reducing loads); intermediate benefit (to transited networks) and end-network benefit (at the target).

Finally, we advise that network providers view this as both a necessary undertaking (cost) to be made efficient, and a profitable service opportunity to be capitalized on. Both demand the same things: efficacy, and cost efficiency as DDoS traffic scales.

Related Research

- Industrial Automation and 5G
- Preparing for Edge Cloud
- Google Cloud Platform – Profile
- Test & Assurance in Cloud Networks
- Open RAN

ABOUT THE AUTHOR

Grant provides a unique combination of management and technical acumen, combined with 30 years of successful innovation in both technology and business models. Prior to Appledore, Grant most recently served in the office of CTO for Ericsson. Through his career, Grant has specialized in transforming telecom software and service businesses in the face of dramatic market and technology shifts, positioning the businesses for survival and growth in new environments.

Grant has deep experience in understanding market and technology shifts, and the consequent business opportunities and threats that these shifts create. Grant's current areas of research focus include closed loop automation, cloud native orchestration & automaton methods, network & cloud security, and the emerging distributed enterprise / edge. He has consistently guided Telcordia and Ericsson software product portfolio to thrive on these changes. Grant holds a Bachelor of sciences from Drew University and a Joint degree (SM-Management, SM-Engineering) from the Massachusetts Institute of Technology.

ABOUT APPLEDORE RESEARCH

Dedicated to the telecom industry, Appledore Research helps its clients navigate rapid change in technology, service innovation and operational practices to positively transform and grow their businesses. We think differently, challenge the status quo and identify the best ways to move forward. We want you to succeed.

Through our expert team of experienced telecom industry analysts and business practitioners, Appledore Research provides you with a unique blend of sector-specific market research reports, strategic advice and marketing support services, either as a part of a subscription package or as individual purchases. We're flexible.

As you'd expect, key focus areas for Appledore's comprehensive research work include Cloud Management, Network Function Virtualization (NFV), Software Defined Networking (SDN), 5G, AI and analytics technologies, service innovation and best operational practice. As a result, our client base spans the complete spectrum of the global telecom industry.

So, we work closely with communication service providers, network providers, software and IT vendors, system integrators, investors and enterprise users. Unbiased and impartial in our views, we are regarded by our clients as a trusted partner. One that's quick to react to disruption in telecom market dynamics and deliver you the most considered and relevant response.

The market research, advice and support, that Appledore brings exclusively to the telecom industry, means better-informed and more efficient business decision-making. It means earlier identification of emerging trends, challenges and opportunities, and fewer mistakes. We're a trusted sounding board, providing that extra bit of objective insight when you need it most.

Insight and analysis for telecom transformation.

@AppledoreVision

Appledore Research

www.appledoreresearch.com

info@appledorerg.com

+1 603 969 2125

44 Summer Street Dover, NH. 03820, USA

Appledore
RESEARCH