NOKIA Bell Labs

The future mode of DDoS attack mitigation

The business case for mitigating volumetric DDoS attacks using Nokia's FP4 silicon

White paper

The need to protect broadband internet infrastructure and its attached customers from distributed denial of service (DDoS) attacks is an increasing concern for network operators.

As DDoS attacks increase in frequency and magnitude, the cumulative cost of the present mode protection solutions based on security appliances is rapidly increasing while their effective traffic coverage is decreasing.

This financial white paper presents the results of a business case analysis by Nokia Bell Labs that compares the present DDoS mitigation solutions with a future mode that leverages peering routers equipped with Nokia's FP4 routing silicon to mitigate volumetric attacks.

Contents

Executive summary

Nokia Bell Labs conducted a business case analysis to model the economics of two modes of distributed denial of service (DDoS) attack mitigation. The study compared the present mode of operation (PMO), which relies on specialized security appliances to detect and remove DDoS traffic from the IP network, with a future mode of operation (FMO): leveraging FP4 silicon to mitigate volumetric DDoS attacks.

The analysis compared the cost of DDoS mitigation for PMO and FMO baseline network models that are each dimensioned for 5 Tb/s, 20 Tb/s, and 50 Tb/s peering capacity. The analysis also calculated the additional cost of peering router capacity, (DDoS) backhaul cost and centralized DDoS scrubbing appliances. The relative cost savings of FMO over PMO are summarized in Figure 1.

Figure 1. Cost savings of FMO using FP4 DDoS filtering versus PMO using scrubbing appliances

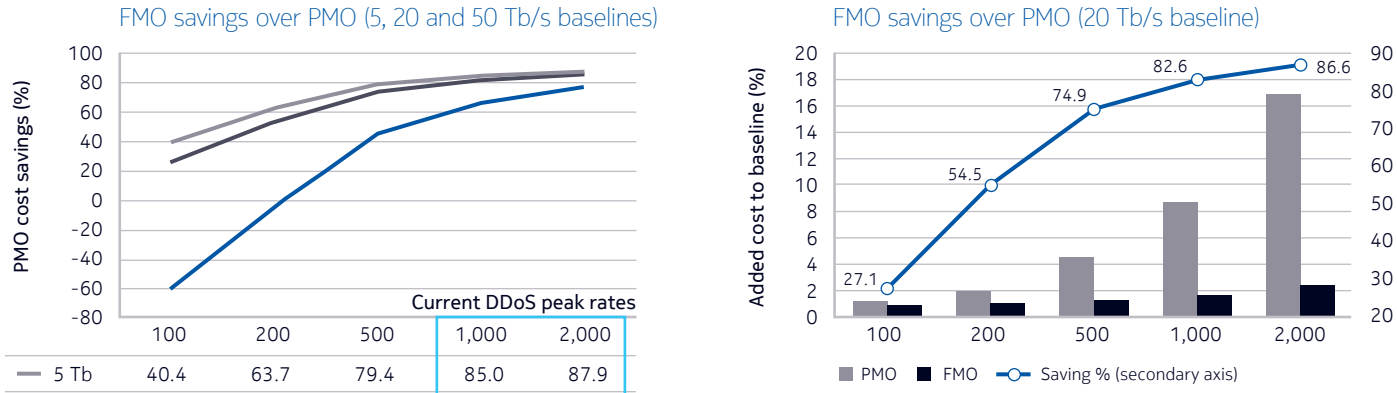

The results show that there is a strong business case for deploying the FMO model. The FMO using FP4 saves between 65 percent and 85 percent on DDoS mitigation cost, depending on network size and assuming current DDoS attack peak rates between 1 and 2 Tb/s.

The DDoS mitigation cost of PMO scales linearly with DDoS peak rates while the FMO approach using FP4 DDoS filtering is virtually unaffected by the peak volume (and frequency) of DDoS traffic.

While the FMO savings achieved with FP4 DDoS filtering are significant, a key advantage lies in FMO's ability to scale with sustainable cost. DDoS attacks keep evolving in magnitude, and the firepower of cloud and IoT allows attackers to keep increasing attack volumes with limited extra costs. The PMO approach based on DDoS scrubbing centers is unable to keep up due to technical scaling limits and incremental costs. Moreover, all PMO countermeasures are reactive in nature because they rely on DDoS detection and peering traffic redirection to intervene in DDoS attacks.

The FMO fully leverages existing network investments because it leverages the FP4 capabilities embedded in the line cards through a simple feature license upgrade. DDoS filters for volumetric reflection and amplification attacks can be proactively applied in-line on all peering interfaces to cost effectively protect the IP infrastructure and its customers from volumetric DDoS attacks.

For more information about the FP4 DDoS mitigation technology, please see the application note "Volumetric DDoS mitigation."

Today's DDoS mitigation challenges

Cyber criminals have made the cloud and the Internet of Things (IoT) powerful and lucrative tools for launching DDoS attacks on essential web services and vulnerable online customers and business assets. Volumetric DDoS attacks are the cyber equivalent of Category 4 and 5 hurricanes, flooding websites and causing internet brownouts, and attacks are becoming increasingly frequent and potent (see Figure 2). In 2017 the number of DDoS attacks increased by 91percent, and 2018 is already setting new records for the largest ever DDoS attack, at 1.7 Tb/s (Memcached)[1].

Figure 2. Evolution of DDoS attacks

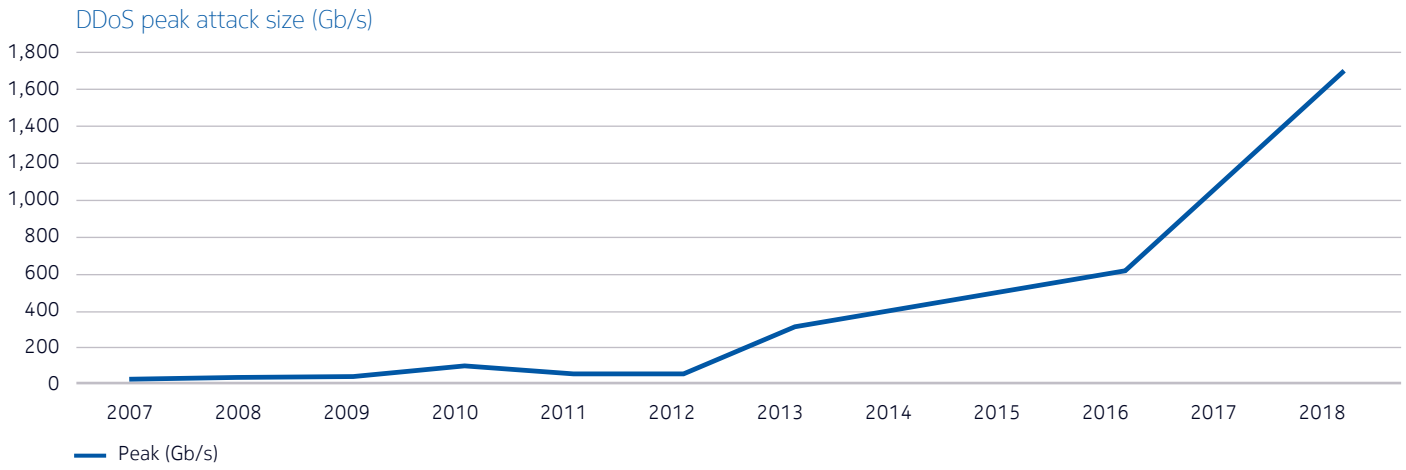

Volumetric DDoS attacks rely heavily on amplification and reflection techniques that exploit vulnerabilities in essential internet service protocols such as the Domain Name System (DNS) and the Network Time Protocol (NTP). For example, when sending a spoofed "ANY request" with the target's IP address to a DNS server, the target will receive server response messages up to 70 times larger than the initial request message. The recent DDoS attacks using Memcached servers were amplified 50,000 times.

A volumetric DDoS attack happens when tens of thousands of hijacked internet devices start sending such requests to the thousands of open DNS and NTP servers that exist on the internet, overwhelming the target with an avalanche of response traffic. By using amplification and reflection techniques, far fewer devices are required to mount volumetric DDoS attacks, and attacking devices are less exposed to the risk of being detected because they use spoofed IP addresses of the target.

The only effective way to block volumetric DDoS attacks is to prevent fake traffic from converging on the target, which can be done by detecting and filtering DDoS traffic in the IP network that connects the victim to the internet. Unfortunately, IP network security has not kept up with evolving demands. Faced with relentless demand for ever more capacity, IP routers have been optimized for moving large traffic volumes at lower cost, which severely limited their capabilities to identify and filter out harmful DDoS traffic. Lacking critical security capabilities in their routers, today's IP networks came to rely on specialized security appliances to detect and remove DDoS traffic on their behalf.

1 Arbor Networks. "NETSCOUT Arbor Confirms 1.7 Tbps DDoS Attack; The Terabit Attack Era Is Upon Us" (blog post). March 5, 2018.

However, security appliances are too costly to deploy in-line on all network traffic, so they are used as a shared network resource in scrubbing centers. When a potential DDoS attack is suspected, traffic for the affected target(s) is rerouted to the scrubbing center to filter out harmful traffic. Precious time is needed to analyze network data and determine the attack vector before the scrubbing center can adequately respond to a threat. State-of-the-art DDoS attacks exploit these vulnerabilities with pulsating, multi-vector attacks that can ramp up in minutes and knock targets and DDoS defenses off-line before they can respond.

A new approach to DDoS mitigation

To address these DDoS mitigation challenges requires a fundamentally new approach that makes the IP network an integral part of the security solution (see Figure 3). Recent advances in routing silicon such as Nokia's FP4 chipset make this possible by combining tremendous forwarding capacity with enhanced packet intelligence and control capabilities. Leveraging highly scalable access control lists (ACLs) and advanced payload filtering capabilities, IP routers can surgically remove volumetric DDoS attack traffic as it enters the network. Moreover, the FMO DDoS mitigation can be deployed in-line on all network traffic, and DDoS ACLs can be configured proactively and permanently to immediately forestall any attacks.

Figure 3. Present and future modes of operation

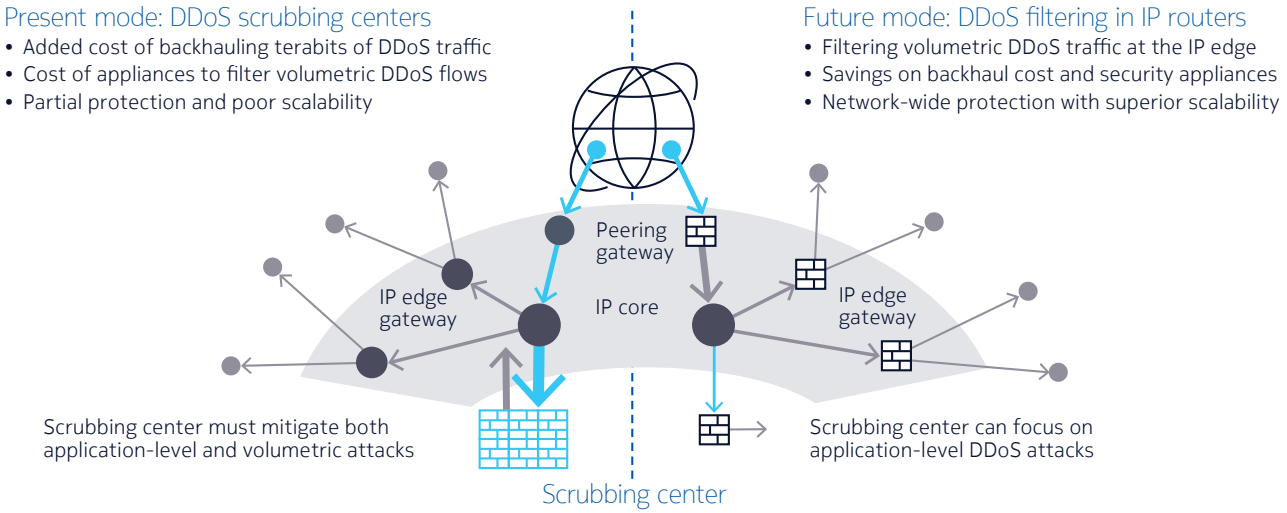

Deploying this technology at network demarcation points such as internet peering routers, data center gateways and subscriber edge gateways establishes a network-wide security perimeter. The IP network now becomes the first line of defense against volumetric DDoS attacks, allowing security appliances to free up precious resources for mitigating session- and application-layer attacks.

This business case analysis by Nokia Bell Labs compares the PMO using scrubbing appliances with an FMO leveraging peering routers equipped with FP4 DDoS filtering.

For more information about Nokia's DDoS protection solutions, please visit the Nokia Deepfield web page.

Assumptions

The business case focuses on incoming DDoS attacks that are launched over the public internet and targeted to connected users who are provided internet access services through the managed network. Although PMO and FMO approaches can also be used to mitigate attacks by on-network users, the clear majority of DDoS attacks originate off-network. To assess the DDoS mitigation cost of the present and future modes of operation, we compared the two network models shown in Figure 4.

Figure 4. PMO and FMO network models

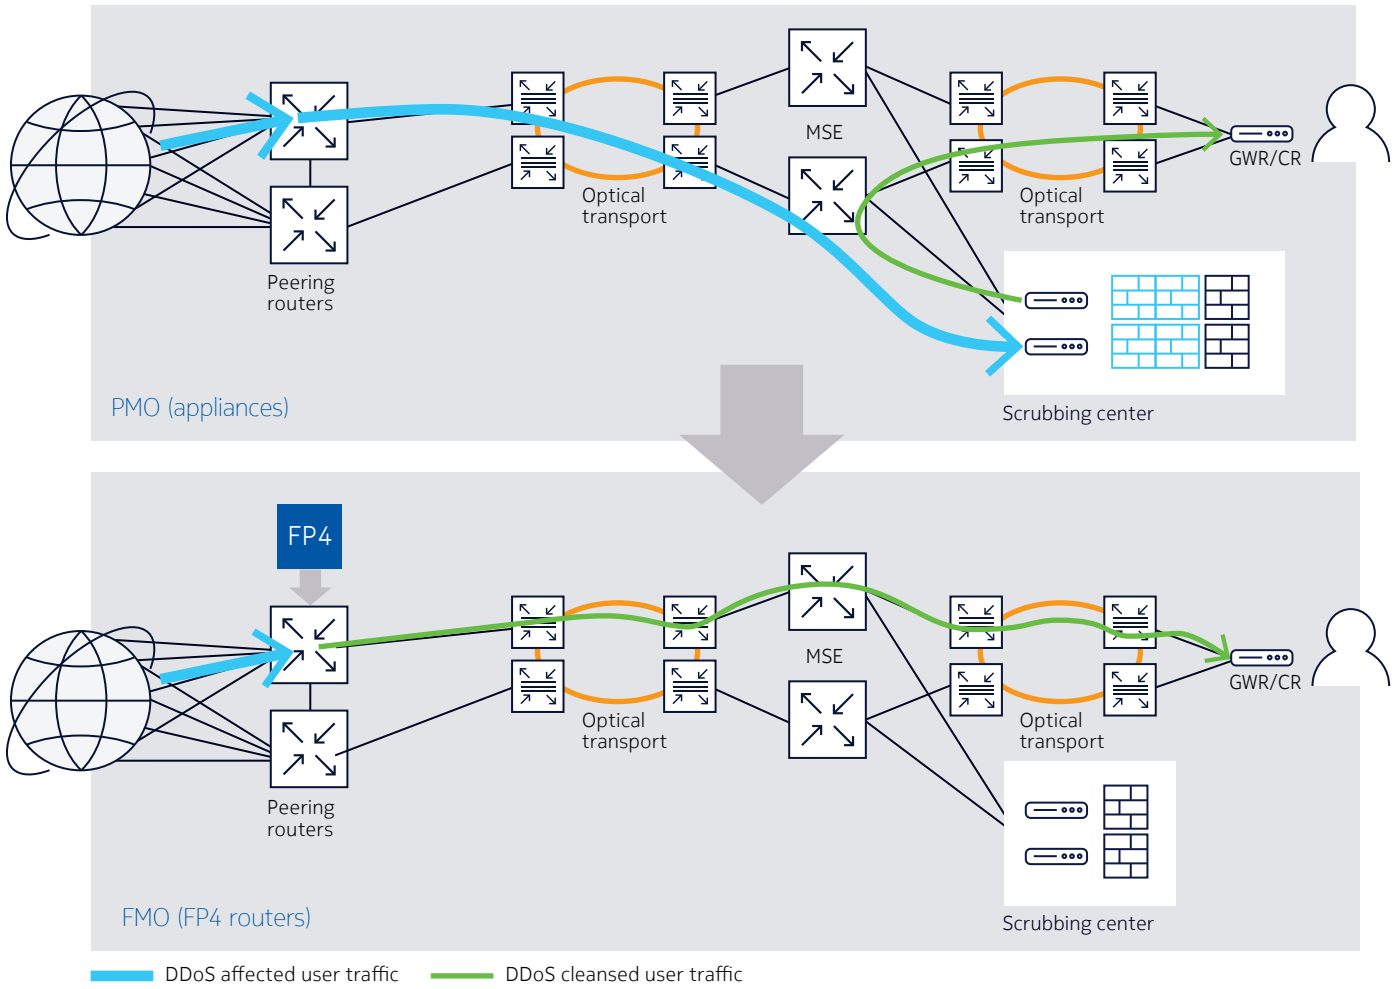

The PMO network model redirects contaminated traffic (blue) to a scrubbing center, which removes DDoS traffic and forwards the cleaned traffic (green) to the user.

The FMO applies in-line DDoS filtering capabilities to remove volumetric DDoS traffic at the peering points. This saves cost on backhaul capacity and frees up scrubbing center resources.

The peering routers are fully redundant, with dual homing of peering links, and are provisioned for a nominal peak load at 50 percent utilization. The optical transport network is also 1+1 redundant.

The hardware configurations and cost assumptions for the peering routers, transport network and scrubbing appliances are identical for the PMO and FMO network models except for the enhanced FP4 DDoS mitigation capabilities, which are enabled only on the peering routers in the FMO.

Table 1 provides an overview of the most important DDoS attack vectors and the frequency of their occurrence as reported by Akamai during Q4 2017. These different types of DDoS attacks listed may leverage different countermeasures depending on the PMO and FMO deployment model.

Table 1. DDoS attack types and countermeasures

DDoS type	Distribution	ALL	PMO	FMO (FP4)	
	%	5-tuple ACL	Appliances	FP4	Appliances
UDP fragmentation	33	100%	0%	0%	0%
DNS reflection	19	0%	100%	100%	0%
NTP reflection	9	0%	100%	100%	0%
CLDAP reflection	10	100%	0%	0%	0%
Direct UDP flooding	7	100%	0%	0%	0%
Chargen reflection	6	100%	0%	0%	0%
TCP (SYN/ACK floods)	7	50%	50%	50%	0%
SSDP reflection	4	100%	0%	0%	0%
SNMP reflection	1	100%	0%	0%	0%
Other/application layer attacks	~5	0%	100%	0%	100%

Basic 5-tuple ACL filtering Stateless payload filtering Various countermeasures L4-7 stateful mitigation

The following DDoS mitigation capabilities are assumed to counter these attacks:

- Peering routers in both PMO and FMO support default 5-tuple ACL capabilities using IP header inspection to:
  - Support DDoS anti-spoofing measures (e.g., Martian address filtering and unicast reverse path filtering)
  - Deny local customer and public internet access to routers and other systems that are part of the managed IP infrastructure
  - Block or rate-limit UDP and TCP ports for internet access that are used by CLDAP, Chargen, SSDP, SNMP (and Memcached) amplification attacks.
- UDP fragmentation attacks are most prevalent (33 percent of all DDoS traffic) but can be mitigated by dropping all first packet fragments below a minimal size and marking all other fragments with a higher drop eligibility than regular packets. This prevents any resource contention in case of a fragmentation attack.
- DNS and NTP attacks are very common (19 percent and 9 percent respectively) and can be amplified to peak rates of well over 500 Gb/s. Mitigation currently requires scrubbing appliances (the PMO), but in the FMO they are mitigated using FP4 payload filters configured on the peering routers.

- High-scale 5-tuple ACL filtering capabilities (i.e., 64,000 to 256,000 entries) are required to mitigate direct UDP flooding attacks because attacks may involve flows from over 100,000 compromised devices. Mitigation is reactive and relies on DDoS detection software such as the Nokia Deepfield solution, which analyzes flow telemetry data to identify and black-list DDoS flows, and white-list trusted flows. Harmful flows are subsequently filtered by peering routers (e.g., using FP4 in the FMO) or by centralized scrubbing appliances in the PMO when lacking sufficient ACL scaling.
- TCP stack attacks, such as SYN/ACK flooding attacks, are not really infrastructure attacks because their aim is to exhaust server resources rather than bandwidth resources. To prevent this from happening, vulnerable targets such as data centers are typically equipped with dedicated in-line firewalls and proxies based on stateful inspection. Although available countermeasures in PMO and FMO can help mitigate a subset of these attacks, there is no clear advantage for either mode in these cases. The study therefore assumes these attacks will be mitigated using dedicated security appliances (assume 2 percent of all DDoS traffic).
- The remainder of attacks are application-layer attacks (Layer 4 to 7) that also require stateful inspection. These attacks are mitigated by specific security devices such as application-layer firewalls, which is a cost for both PMO and FMO (assume 3 percent of all DDoS traffic).

Modeling the cost of DDoS traffic is quite different from modeling regular user traffic, and some important observations and assumptions were made that influence the outcome:

- Volumetric DDoS traffic sizing and modeling must assume worst-case conditions:
  - DDoS attacks can occur at any time but are most effective—and more likely—at peak hours.
  - There is no correlation between DDoS attack peak rates and network size, but larger networks will typically see more attacks because they present a larger attack surface (more targets).
  - Sufficient mitigation resources must be available to handle the largest DDoS attack peaks.
  - DDoS traffic can have any payload (randomized) or composition (single-/multi-vector attack).
- When detecting a DDoS attack, it's unknown what proportion is fake traffic.
  - DDoS traffic is only a subset of all traffic addressed to a target, but to filter out fake traffic it is necessary to inspect and scrub all traffic destined for the target.
  - For volumetric DDoS attacks to be successful, the proportion of fake traffic must typically outweigh the amount of real traffic, assuming 50 percent capacity utilization in normal conditions.
  - DDoS attacks might not be focused on one single target IP, but instead target one or multiple subnets. These are known as "infrastructure attacks" and are harder to deal with by the PMO because DDoS traffic is scattered but will still create congestion further downstream.
- The study focused on only DDoS mitigation cost. It assumes that the DDoS detection capabilities available to PMO and FMO are equal in cost and are infallible (no false negatives or missed positives).

Study details

The study compared three network models with 5, 20 and 50 Tb/s peering capacity, and compares the cost to mitigate DDoS attacks with peak rates of 100 Gb/s, 200 Gb/s, 500 Gb/s, 1 Tb/s and 2 Tb/s.

Tier 1 network with 5 Tb/s peering capacity

The cost comparison is done in terms of added percentage and dollar amount over the network baseline cost of a network dimensioned without DDoS traffic. Figure 5 compares the cost of protecting PMO and FMO network models dimensioned for 5 Tb/s peering capacity.

Figure 5. DDoS cost comparison for a 5 Tb/s network

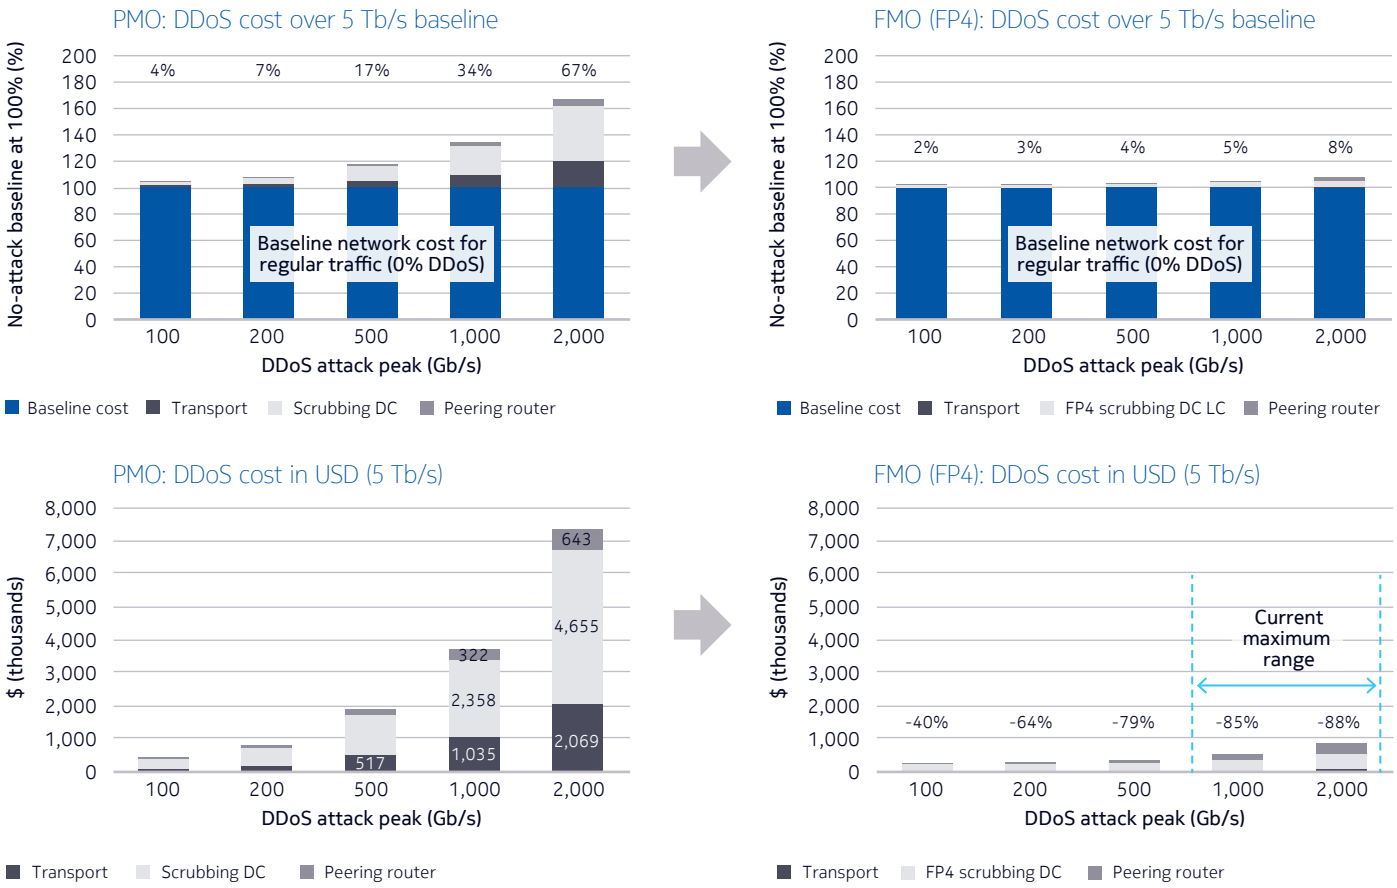

The DDoS mitigation cost represents the additional cost in peering routing capacity, backhaul capacity and scrubbing resources required to mitigate attacks ranging from 100 Gb/s to 2 Tb/s.

- The PMO cost increases from +4 percent for 100 Gb/s attacks to +67 percent for 2 Tb/s attack peaks.
- The FMO cost increase is more moderate, growing from +2 percent for 100 Gb/s to +8 percent for 2 Tb/s attack peaks, to account for 5 percent of all DDoS traffic being handled by dedicated appliances.
- The **FMO using FP4 saves approximately 85 percent on the PMO cost**, considering current volumetric DDoS peak rates ranging between 500 Gb/s and 1 Tb/s. The PMO cost savings are achieved on customer-facing peering capacity, DDoS backhaul cost and DDoS scrubbing appliances.

Tier 2 network with 20 Tb/s peering capacity

Figure 6 compares the cost of protecting PMO and FMO network models dimensioned for 20 Tb/s peering capacity.

Figure 6. DDoS cost comparison for a 20 Tb/s network

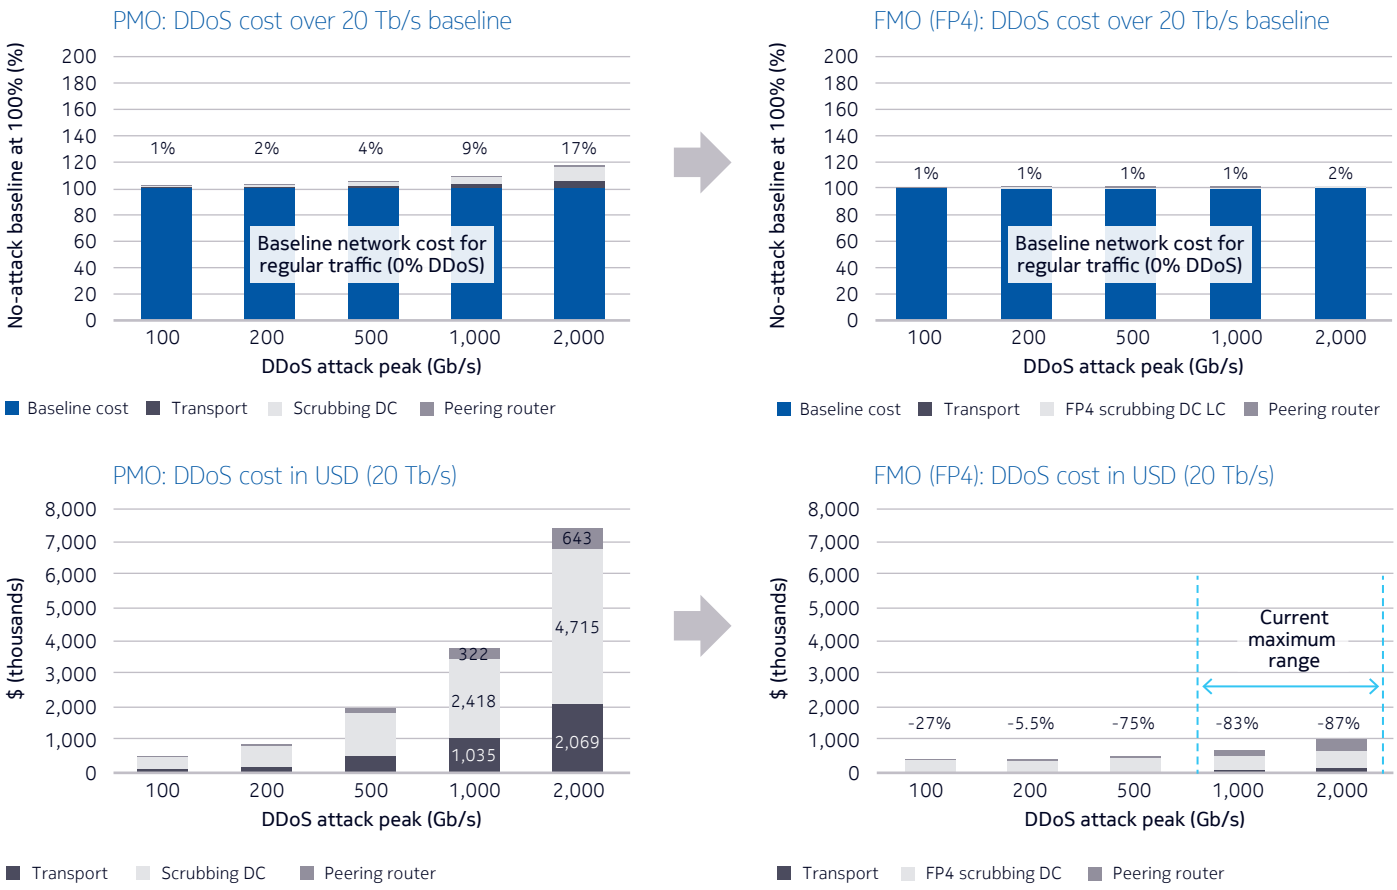

- The PMO cost increases from +1% for 100 Gb/s attacks to +17 percent for 2 Tb/s attack peaks, but traffic coverage is declining as well, from 20 percent to 5 percent of all network traffic at 1 Tb/s peak rates.
- The FMO cost increases from +1 percent for 100 Gb/s to +2 percent for 2 Tb/s peaks, but covers 100 percent of network traffic.
- The **FMO using FP4 saves approximately 85 percent on the PMO DDoS mitigation cost**.

Tier 3 network with 50 Tb/s peering capacity

Figure 7 compares the cost of protecting PMO and FMO network models dimensioned for 50 Tb/s peering capacity.

Figure 7. DDoS cost comparison for a 50 Tb/s network

- The PMO cost increases from +1 percent for 100 Gb/s attacks to +7 percent for 2 Tb/s attack peaks, but traffic coverage is declining to only 2 percent of all network traffic assuming 1 Tb/s peak rates.
- The FMO cost is +1 percent over baseline cost and is virtually unaffected by DDoS attack rates.
- The **FMO using FP4 saves approximately 65 percent to 78 percent on the PMO** at current DDoS peak rates.

NOKIA Bell Labs

References and resources

1. Netscout/Arbor Networks. Arbor confirms 1.7 Tbps DDoS attack. The terabit era is upon us (blog post). March 5, 2018.
2. Nokia Networks. Deepfield solution web page.
3. Nokia Networks. FP4 web page.
4. Nokia Network. "Volumetric DDoS mitigation: Making the network part of the solution" (application note). October 2017.

Abbreviations

ACK	Acknowledge
ACL	access control list
CLDAP	Connection-less Lightweight X.500 Directory Access Protocol
DDoS	distributed denial of service
DNS	Domain Name System
FMO	future mode of operation
IoT	Internet of Things
MSE	Multi-Service Edge
NTP	Network Time Protocol
PMO	present mode of operation
SNMP	Simple Network Management Protocol
SSDP	Simple Service Discovery Protocol
SYN	Synchronize
TCP	Transport Control Protocol
UDP	User Datagram Protocol

About Nokia Bell Labs

Nokia Bell Labs is the world-renowned industrial research arm of Nokia. Over its 90-year history, Bell Labs has invented many of the foundational technologies that underpin information and communications networks and all digital devices and systems. This research has resulted in 8 Nobel Prizes, two Turing Awards, three Japan Prizes, a plethora of National Medals of Science and Engineering, as well as an Oscar, two Grammys and an Emmy award for technical innovation. Nokia Bell Labs continues to conduct disruptive research focused on solving the challenges of the new digital era, defined by the contextual connection and interaction of everything and everyone, as described in the book, The Future X Network: A Bell Labs Perspective. www.bell-labs.com

Nokia Oyj
Karaportti 3
FI-02610 Espoo, Finland
Tel. +358 (0) 10 44 88 000

Document code: SR1906036268EN (June) CID202100