Leveraging Software-Defined Networking for Security Intelligence against Distributed Denial of Service Attacks in the Data Centre

Ivan Iordanov

***Abstract*—**Distributed Denial of Service (DDOS) attacks are particularly damaging for commercial Data Centres (DC) – they affect the operator, their customers and the end users. The high capacity upstream connectivity of those organisations drives up the price of purely volumetric attacks against them. This makes other flavours of DDOS, which can be protected from within an organisation, the preferred tools of attackers. In order to mount a proper defence, however, timely actionable intelligence is needed. Collecting said Security Intelligence (SI) can be accomplished by means of an Intrusion Detection System (IDS). Finding patterns of multi-day subtle sophisticated attacks, or performing post-mortem analysis of suffered breaches, requires keeping huge amount of data backlog. This study presents an overview of DDOS attacks from the point of view of actors involved and enumerates various ways to get around the storage requirements by using flexible traffic selection mechanisms. A formula for calculating needed storage capacity is provided. The resultant recommendations are particularly applicable for Software-Defined Networking (SDN) environment where dynamic rollout of new interpretative rules can be leveraged in order to perform highly specialised operations on multi-purpose hardware.

***Keywords*—** Distributed Denial of Service (DDOS), Intrusion Detection System (IDS), Security Intelligence (SI), Software-Defined Networking (SDN)

I. Introduction

DATA centres (DC) are target rich environment for distributed denial of service (DDOS) attacks. Nevertheless, not all attack vectors originate from outside. The concentration of servers administered by diverse groups adhering to varying standards of information security facilitates compromising at least some machines. Depending on the success of the operation, they can be included in botnets, be taken over and used to perpetrate an information extraction under the cover of DDOS or even, especially in commercial settings, be legitimately leased by the attacking party and 'accidently' left with exploitable backdoors for the purposes of plausible deniability. Another concern is the saturation of servers with high speed uplinks that can become vectors in amplification DDOS targeting third parties. All of the above, leaves the DC owners and/or operators open to tort liability [1].

Ivan Iordanov is with the Department of Computing Systems of Faculty of Mathematics and Informatics, University of Sofia, 1164 Sofia, Bulgaria (e-mail: iiordanov@ucc.uni-sofia.bg).

II. Anatomy of a DDOS Attack

A denial of service (DOS) attack represents an attempt to overwhelm the resource capabilities (be it computing or bandwidth) of a given group of network or termination devices so that the services provided to legitimate users of the former experience either degradation or total outage. This is achieved by generating illegitimate traffic and requests towards the attacked device(s). It should be noted that not everything that looks like a DOS is necessarily a DOS. For example, in the 90ies the Internet consisted of many weak servers and slow connection. As such, hostmasters were often victims of their own success – a launch of a new piece of software presented on the website and/or citation at a popular aggregator site would result in upsurge of interest effectively bringing down the operation for hours and even days while new capacity was being installed.

A DDOS is just DOS attack originating from multiple sources. Those sources are usually part of one or more botnets – a collection of computing devices unwittingly taken over by the attacker beforehand (often with the help of a Trojan). DDOS has several significant advantage over its predecessor, which explains its omnipresence:

- The source is harder to determine;
- The pooling of resources makes overwhelming the target cheaper;
- The attacker can orchestrate the operation without directly engaging the target.

Another ingenious type of DDOS is the amplification reflection DDOS (also known as Distributed Reflection DOS or DRDOS). This attack utilises servers providing UDP based services (such as NTP or DNS) as force multipliers – a spoofed request is sent to them with the IP address of the target as source. The resulting reply is 5 to 500 times bigger in size than the original request [2].

After this brief summary, several distinct actors in the perpetration of attacks can be differentiated.

A. Instigator

The instigator is the one orchestrating the attack. In order to hide his hand, he can use various intermediate points connected with a myriad of crypto tunnels before contacting the botnets. This serves the dual purpose of obscuring his

identity and/or affiliation and protecting the command and control (C&C) channels from disruption [3].

B. Attack vector

The attack vector is a distinct grouping of devices generating malicious traffic. A DDOS can have multiple attack vectors which together form the totality of the attack. They are directed via the instigator's C&C and can be rendered inert by disrupting the later.

While in a classic DOS the instigator and attack vector are one and the same, this is not the case in DDOS. Typically, hosts are infected beforehand by a Trojan and formed into a botnet. There are, however, instances in which distributed application infrastructures such as various peer-to-peer networks can be subverted for similar purpose [4].

C. Force multipliers

Force multipliers are unique to amplification attacks. They represent machines running UDP based services for which exists a query whose size is significantly smaller than the response. A request with spoofed source IP address is sent from an attack vector in order to achieve the amplification effect.

The multiplication factor can be further increased by taking control (legally – provisioning; or illegally – hacking) of servers that interact with the force multipliers. An example would be creating a TXT record on a controlled DNS server and querying open recursive DNS servers for it [5].

If the amplification factor is less than 1 (ping can be considered as etalon here – ICMP ECHO reply has the same payload as the ICMP ECHO request which provoked it [6]), the force multiplier is used only to obscure the origin of the attack.

D. Target

The target of the attack is its intended victim. This can be either a whole network – volumetric attack aimed at exhausting the available upstream bandwidth, or a particular server.

In the latter case, preventing the traffic from reaching the target (even if it has already breached the victim's autonomous system) defeats the attack since the server's computing resources are not impeded. In the former, however, even if all attack traffic is dropped on the upstream port, legitimate traffic need to contest the last mile link with the attack flood, thus resulting in service outage or degradation.

III. REASONS FOR DDOS ATTACKS

Year-on-year, the volume and size of DDOS attacks has been consistently increasing since the last decade of the previous century [7]. There are multiple drivers behind this trend – increase in the number of traditional terminating devices, evolution of Internet of Things (IoT), lowering of the entry knowledge threshold due to proliferation of attack tools and techniques, emergence of divergent motives. Even though this paper does not aim to discuss solutions for the DDOS phenomenon by tackling its primary founts, it would behove us to examine each of the four drivers enumerated above in more details.

The rapid growth of the Internet population (400 fold increase of connected users between 1994 and 2014 [8]) observed in the last decades is important for several reasons:

- The shear rapidity of the process ensured widespread lack of cyber security culture. Thus the potency of social engineering attacks is heightened.
- Each user possesses multiple on-line presences. The majority (or totality) of which are password protected. Security key overlap (between different people and/or identities) is inevitable. Thus the potency of dictionary attacks is heightened.
- Most people have multiple fully versatile computational devices (smartphones, tablets, work/home computers/laptops, etc.) which are often pre-configured to connect to each other with minimal hassle. This means that breaching the security on one can provide an easy path to the rest. Additionally, despite all attempts at raising awareness and automating the process, regular installation of software security patches is not common enough – there are many devices unsecured against known exploits months and even years after the discovery of the same [9].

The proliferation of IoT (while only 40.4% of the world population in 2014 is on the grid [8], the number of devices tops 20 billion [10]) has somewhat different aspects with regards to DDOS. The software of IoT devices is generally firmware with limited capabilities for communication, security and upgrade. Additionally, people do not always actively use the IoT features. This makes IoT popular as both vector and target of DDOS:

- Many devices are left with default configuration and/or passwords.
- There are much weaker automated processes for patching up exploits on their firmware compared to a fully featured operating system (OS).
- By definition the devices are capable of communication information regarding their operation in solicited or unsolicited manner.
- IoT device have low processing power and memory compared to traditional computers and, as such, are easier to DDOS than the later.

Another challenge to the security community is that the vast majority of hackers have become tool users – possessing passing familiarity with programming, networking and system architecture, but not capable of executing an attack using only their own code. Even the people creating the tools utilised by both white and black hats are overly specialised. Indeed, comparing the hackers of the 1990s to those of today is like contrasting the medieval blacksmiths with the compartmentalised steel industry at the end of the 19th century.

Finally, the last two decades saw a profusion of new

motivations. While at the dawn of the Internet age curiosity and general love of mayhem were the main instigators of attacks, as the technology, with its perfect deniability, matured and the cyber economy grew in size and importance, more traditional factors started to weigh in:

- Continuation of politics with other means (i.e. cyberwar [11,12] and cyber-espionage [13])
- Promotion of political agendas by state and non-state actors (i.e. hacktivism [14, 15])
- Transformation of traditional criminal activities such as extortion (i.e. DDOS ransom demands [7])
- Cover for other illegal, clandestine or terrorist activities (i.e. pre-operational surveillance, data extraction, malicious code injection, inciting market panic [16], etc.)

IV. SI COLLECTION MANAGEMENT IN SDN ENVIRONMENT

A. SI Collection

Security Intelligence (SI) aims to provide to the Information Technology (IT) sector, some of what traditional Intelligence Agencies have been serving to state actors for centuries – namely actionable information that can be used to mitigate threads against the organisation, as well as analysis on trends that can be exploited against less well-informed commercial rivals. In IT, the crucial intelligence is traffic patterns – volume, destination, source, protocol, etc. Being able to distinguish the critical patterns of attack from the noise of the legitimate traffic is as important to securing the present as identifying the outliers is to preparing for future threads.

Let us consider the different types of information of interest in a commercial DC:

- What traffic is being sent to unallocated IPs of the DC's autonomous system (AS)? – This clearly delineated illegitimate traffic. It can have three sources – victims of DDOS replying to spoofed messages (potentially, data on unseen attacks can be collected that way); botnets performing ping-sweeps and port mapping (i.e. pre-operational surveillance); or researchers and otherwise benign actors doing the same (ironically, it is more likely that they will do this from a single IP address).
- What traffic is being sent to the DC servers? – This is the critical data needed to recognise and prevent DDOS. It includes both the customer and attack traffic. Information regarding what the servers are actually sending can be used to distinguish the legitimate flows.
- What traffic is originating from the servers? – This most likely describes legitimate traffic and can thus be used to narrow down the suspicious outside traffic – healthy TCP connections would exchange traffic past the SYN stage, ICMP echo or DNS replies should only be received after an outgoing request has been detected, if a UDP packet has originated from the server, then its reply would be legitimate, as well, etc. Additionally, potentially suspicious behaviour of a server indicating that the machine is compromised might be detected.

B. SI Retention

One should not forget that SI is designed to be more than an Intrusion Detection System (IDS). Indeed, it is envisioned as a holistic union of all aspect of IT security. This means that logging management, event monitoring and network forensics are also considered to be part of it. It is this last aspect in particular that requires retention of old information.

Thus it is appropriate to discuss the benefits of forensics. Much like in criminology, network forensics are used to reconstruct events of interest in the past. The aim is to obtain information on the modus operandi (MO) of malefactors, ascertain their affiliation and identity, formulate new preventive measures (which usually equates to updating the attack fingerprints in order to allow better early detection) and, ultimately, collate admissible evidence for court proceedings (however unlikely it is that this will amount to anything [17]).

The question remains – what is pertinent information? As a rule, DDOS attacks originate from outside the AS. This means that OSI layer 3 (IPv4/IPv6) is the lowest level of detail that is of interest. There remains the outlier possibility that in a commercial DC environment a server has been compromised and is being used to attack other machines in the network. However, prudent design considerations suggest that application servers of different customers should not be sharing a common layer 2 broadcast domain (easily accomplished with the use of isolated ports, even if on the same vlan).

The IPv4 header can be between 20 and 60 bytes [18] and the same for TCP [19], while the header for ICMP and UDP is 8 bytes [6, 20]. An IPv6 header can extend from a minimum of 40 bytes to the maximum path MTU [21]; considering it will be traveling through the commercial Internet, MTU is unlikely to exceed 1500 bytes, though if the vector is internal to the DC, a jumbo MTU of 9600 bytes might be in effect. A summary of the above can be found in Table I.

TABLE I. MIN/MAX HEADER SIZE, BYTES

	Pure L3 header	**w/ TCP**	**w/ UDP**	**w/ ICMP**
IPv4	20/60	40/120	28/68	28/68
IPv6	40/1500	60/1500	48/1500	48/1500

Since an SI collection system is more than a DDOS IDS, or even a general purpose IDS, a measure of forensic capabilities are implied and desired. To this end, at least part of the packet payload should also be stored for analysis – all OSI layer 7 exploits rely on particular strings being carried therein. Naturally, the question of what precisely should be stored and for how long depends on a cost benefit analysis between the cost of storage and the expected gain in security framed by any existing, or expected impending future, contractual obligations with DC customers and/or government regulations.

1) Retention – The retention policy must ensure that historical data is present in enough detail to reconstruct an event of interest days after its occurrence.

a) *24 hours* – DDOS attacks that take days are rare outliers. In any case, response time for an organisation heavily reliant on Internet connectivity for its business model (thus being particularly susceptible to cyber-attacks) is less than 60 minutes. Most attacks exhaust themselves in hours. This means that a day of data retention will be sufficient to understand the event itself, provided suitably qualified personnel are available 24/7 to perform analysis.

b) *72 hours* – Retention of 3 days backlog is convenient when there are only a few people able to analyse patterns and they all work on standard business hours. Additionally, this covers not only the attack itself, but also whatever pre-operational surveillance has been performed in preparation (ping sweeps, port probes, exploit vulnerability checks, etc.).

c) *168 hours* – A full week of traffic information is only really needed for deep forensic analysis of highly sophisticated attacks used as cover for information extraction or system penetration acts.

2) Granularity – How much information out of each packet is stored for analysis has important implications for determining the storage capacity requirements for a given retention policy.

a) *Pure metadata* – Having IP packet headers is the minimum information necessary in order to generate fingerprints. This is enough forensic data to re-create most DDOS type attacks, but does not contribute to understanding new exploits or more sophisticated types of attack aimed at damaging and/or extracting information from IT devices.

b) *Partial payload* – This level of granularity collects the OSI layer 3 headers, as well as part of the packet payload. The aim is to better understand the nature of the perpetrated attack. For example, the beginning of a UDP packet coming from a DNS server (i.e. part of reflection-amplification DDOS) can identify the resources being polled on the attack vector device. This in turn can be used to determine servers which have been compromised (i.e. hackers have uploaded abnormally large dummy records), are under the attackers control (i.e. a zone with only a couple of A records, but one or more excessive TXT records) or are suffering from systemic vulnerability. It should be noted that, the attacker can hide the malicious commands in the second half of a packet (if we consider a known web server vulnerability like SQL injection, the actual commands are preceded and/or interspersed with innocuous requests whose presence is expected by the application server).

c) *Total capture* – Another way to go about solving the problem of adversaries occluding relevant information amidst seemingly trivial communication is to capture everything and perform deep analysis on part of the traffic – either during post-mortem following an attack or on regular basis based on a set of pre-selection criteria. A great benefit of this approach is that new patterns can be learned by re-evaluating traffic which has been initially judged innocuous, but was, subsequently, found to coincide with an attack and is thus retroactively relabelled as suspicious.

3) Scope – What traffic is captured directly impacts the storage requirements. Being too restrictive can also impair the learning capabilities of the IDS system.

a) *Bogon traffic* – Adopted from hacker parlance, the term denotes illegitimate traffic whose forged source IP address belongs to unallocated or special (private, experimental, reserved) address space. Best practices suggest to drop such packets directly on the Internet border router (BR). However, since they are clearly traffic that is not part of regular communication, intercepting them (most expediently by means of policy based routing (PBR)) for analysis is prudent. The downside is that those packet are mostly part of reconnaissance and volumetric attacks.

b) *Attack traffic* – When an Intrusion Detection System/Intrusion Prevention System (IDS/IPS) finds a match against known attack fingerprint, the packets conforming to the pattern are slated for discard. This action is most fruitions if performed on the closest point to the source (i.e. the BR). Much like with the previous case, one can instead redirect it to an IDS collector via PBR for further analysis. The benefits are improving the quality of existing fingerprints so as to reduce the percentage of false positives, as well as documenting new flavours of known attacks (i.e. discovering unknown vector for amplification-redirection DDOS).

c) *Suspected traffic* – This method presumes that an IDS can detect at least some of the attack traffic towards a target. When this happens, all traffic towards the target is redirected to the collector for further analysis allowing the discovery of new fingerprints.

d) *All traffic* – Collecting all traffic gives the best forensic and analysis options, but elevates the storage requirements.

4) Sampling – Not subjecting every packet to examination and retention can alleviate storage and privacy concerns.

a) *Deterministic sampling* – A method favoured by companies managing the estates of multiple customers that can be likened to mosaic intelligence gathering. It consist of taking every one in n packets. If the solution is applied in a single site, n should be small (between 2 and 1000). If it is used to collect information from diverse dispersed locations and/or jurisdictions a higher number is preferable (1 000-10 000).

b) *Random sampling* – Every n packets one is taken for analysis. Similar to deterministic sampling but provides better spread for higher values of n.

c) *Sequence sampling* – Variation of random sampling where a sequence of k packets with the same source/destination addresses is sampled out of every n. It can be likened to sequence sampling in animal behavioural science where a particular interaction is the focus of observation instead of any of its actors [22]. It improves attack detection for multifaceted and/or multistage attacks since a single packet can easily be misjudged as innocuous without context (i.e. TCP SYN as part of normal protocol operation, as opposed to SYN flood).

d) *Total sampling* – Collecting all traffic gives the best forensic and analytical options, but elevates the storage requirements.

5) *Capacity* – A consistent 24-hours 10Gbps traffic flow would fill up 108TB of storage space (equal to, for example, 36×3TB disks). In any case, a minimum of RAID0 2-disk array must be used, if employing SATA III interfaces (6Gbps), to accommodate each 10Gbps of line speed. When calculating storage requirements, either a historical daily average should be used or a rule of thumb such as 60 (or 70, depending on company policy) percent average utilisation for the upstreams and an 8% risk budget (2 hours of maximum utilisation for every 24 in order to simulate a volumetric DDOS attack).

C. Storage selection

The above analysis leaves us with five variables to take into consideration when evaluating how much storage in needed to accommodate the SI needs of an organisation. Two of them (namely retention and capacity), lend themselves to exact computation. The other three are more difficult to evaluate with absolute certainty. A comparison for the relative storage requirements of each option can be found in Table II.

TABLE II. STORAGE REQUIREMENTS FOR DIFFERENT TRAFFIC INSPECTION POLICIES

	Low	Medium	High
Granularity	(a) Pure metadata	(b) Partial payload	(c) Total capture
Scope	(a) Bogon traffic	(b) Attack traffic (c) Suspected traffic	(d) All traffic
Sampling	(a) Deterministic (b) Random	(c) Sequence	(d) Total

Ultimately, the size of the necessary storage D, in terabytes, for a given set of SI policies can be calculated by the equation

$$D = 0.45 \times T \times B \times (u + r) \times n \times p \times m \,, \quad (1)$$

where *0.45* is coefficient for converting Gbps in TB per hour, T is chosen retention policy in hours, B is total upstream bandwidth in gigabits per second, u is upstream utilisation, r is risk budget, n is percentage of captured traffic based on granularity policy, p is percentage of captured traffic based on scope and m is the same due to sampling policy.

D. SDN Environment

The remarkable thing about Software-Defined Networking (SDN) is the ability to dynamically rollout not only new configurations, but also new interpretative rules. In effect, this means that new or changed standards and protocols that had traditionally required hardware and/or firmware upgrade can be distributed as standalone applications without disrupting the ongoing operation of the affected nodes.

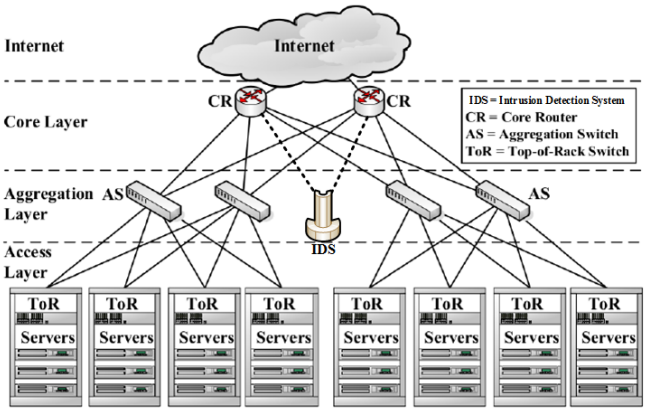

Fig.1 Classical IDS deployment in SDN and non-SDN environments

It also allows for conditional SI collection as described in sections IV-B-2-b, IV-B-3-b, IV-B-3-c, IV-B-4-a, IV-B-4-b or IV-B-4-c to be applied on general purpose hardware located anywhere in the network (edge, aggregation or access). Additionally, points of traffic redirection can vary based on link utilisation (i.e. a classical traffic capture performed on the edge, as show on Fig. 1, can instead be moved to the aggregation layer, as show on Fig. 2, even on the fly in case of link saturation on the core routers). In essence, all middle ground policies become viable alternatives without necessitating rollout of specialised equipment. Indeed, their better precision makes the preferred solution.

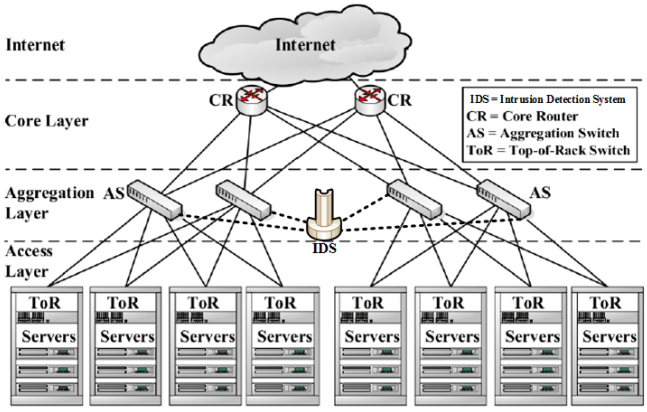

Fig. 2 Possible IDS deployment in SDN environment

V. CONCLUSION

Data centres are saturated with servers – each and every one of them a target interest for at least some attackers. Protecting this estate from penetration is vital for the commercial

prospects of the organisation, and its customers, and yet, often, there is no end-to-end security – the ability of provider to enforce regular updates for most, or even any, of the applications on customer managed devices is sketchy at best, even if enshrined in contract.

On the other hand, the high capacity upstream connectivity provides protection on its own from most volumetric attacks. Of course, with record DDOS attacks reaching 500Gbps [23] and 600Gbps [24] anyone apart from Tier 1 providers can be overloaded. The DCs are also targeted by other types of DDOS – those that can be protected against within an organisation. The deployment of an SI system with a strong IDS component can be most be beneficial for repulsing them.

The strong trend of SVN adoption in DC [23], will allow organisation to deploy SI using the more versatile and discerning polices described in this study.

REFERENCES

[1] A. P. Karanasiou, "The changing face of protests in the digital age: on occupying cyberspace and Distributed-Denial-of-Services (DDoS) attacks," International Review of Law, Computers & Technology, vol. 28, no. 1, pp. 98-113, 2014.

[2] C. Rossow, "Amplification Hell: Revisiting Network Protocols for DDoS Abuse", 2014 Network and Distributed System Security (NDSS) Symposium on 2014 February 23-26, San Diego, CL, USA.

[3] L. Watkins, K. Silberberg, J. A. Morales and W. H. Robinson, "Using inherent command and control vulnerabilities to halt DDoS attacks," 2015 10th International Conference on Malicious and Unwanted Software (MALWARE) on 2015 October 20-22, Fajardo, PR, USA, pp. 3-10, IEEE, 2015.

[4] H. Koo, Y. Lee, K. Kim, B. H. Roh and C. Lee, "A DDoS attack by flooding normal control messages in Kad P2P networks", Advanced Communication Technology (ICACT), 2012 14th International Conference on 2012 February 19, Pyeong Chang, South Korea, pp. 213-216, IEEE, 2012.

[5] M. Anagnostopoulos, G. Kambourakis, P. Kopanos, G. Louloudakis and S. Gritzalis, "DNS amplification attack revisited", Computers & Security, vol. 2013, no. 39, pp.475-485.

[6] J. Postel, "Internet Control Message Protocol - DARPA Internet Program Protocol Specification," RFC 792, USC/Information Sciences Institute, September 1981.

[7] S. Mansfield-Devine, "The growth and evolution of DDoS", Network Security, vol. 2015, no. 10, pp. 13-20, 2015.

[8] "Number of Internet Users (2016) - Internet Live Stats", 2016, [Online], Available: http://www.internetlivestats.com/internet-users/

[9] A. Bartholemy and W. Chen, "An examination of distributed denial of service attacks," InElectro/Information Technology (EIT), 2015 IEEE International Conference on 2015 May 21, Dekalb, IL, USA, pp. 274-279.

[10] D. Evans, "The Internet of Things: How the Next Evolution of the Internet Is Changing Everything", Cisco Internet Business Solutions Group (IBSG), April 2011.

[11] M. Lesk, "The new front line: Estonia under cyberassault", IEEE Security & Privacy, vol. 2007, no. 4, pp. 76-79, 2007.

[12] S. Collins and S. McCombie, "Stuxnet: the emergence of a new cyber weapon and its implications", Journal of Policing, Intelligence and Counter Terrorism, vol. 7, no. 1, pp. 80-91, 2012.

[13] G. O'Hara, "Cyber-Espionage: A growing threat to the American economy", CommLaw Conspectus, vol. 19, p.241, 2010.

[14] G. Lucas, "Ethical Challenges of 'Disruptive Innovation': State Sponsored Hactivism and 'Soft'War," in Evolution of Cyber Technologies and Operations to 2035, Springer International Publishing, 2015, pp. 175-184.

[15] U. Tatar and M. Celik, "Hacktivism as an emerging cyberthreat," in Terrorism Online: Politics, Law and Technology (2015), Routledge, 2015, pp. 54-71.

[16] I. Kilovaty, "Rethinking the Prohibition on the Use of Force in the Light of Economic Cyber Warfare: Towards a Broader Scope of Article 2 (4) of the UN Charter," JL & Cyber Warfare, vol. 2014, no. 4, p. 210, 2014.

[17] J. O. Nehinbe, "A comparative study of attributes for gathering admissible evidence in the investigation of distributed denial of service (DDoS) attacks," International Journal of Internet Technology and Secured Transactions, vol. 4, no. 2-3, pp. 121-138, 2012.

[18] J. Postel, "Internet Protocol - DARPA Internet Program Protocol Specification," RFC 791, USC/Information Sciences Institute, September 1981.

[19] J. Postel, "Transmission Control Protocol - DARPA Internet Program Protocol Specification," RFC 793, USC/Information Sciences Institute, September 1981.

[20] J. Postel, "User Datagram Protocol," RFC 768, USC/Information Sciences Institute, August 1981.

[21] F. Gont, V. Manral and R. Bonica, "Implications of Oversized IPv6 Header Chains," RFC 7112, Internet Engineering Task Force (IETF), January 2014.

[22] J. Altmann, "Observational study of behavior: sampling methods," Behaviour, vol. 49, no. 3, pp.227-266, 1974.

[23] "Arbor Networks - Worldwide Infrastructure Security Report 2016", vol. 11, p. 12, Arbor Networks, 2016, [Online], Available: https://www.arbornetworks.com/images/documents/WISR2016_EN_Web.pdf

[24] S. Khandelwal, "602 Gbps! This May Have Been the Largest DDoS Attack in History", 2016, [Online], Available: http://thehackernews.com/2016/01/biggest-ddos-attack.html