NEXUSGUARD®

Threat Report

Distributed Denial of Service (DDoS)

Q3 2018

Contents

Key Observations 02

Quarterly Focus

DDoS Attacks the “Mongol Way” 03

Junk Traffic Hijacking Legitimate Traffic 04

Implications for ASN-level CSPs 06

Attack Vectors and Their Targets 07

DDoS Activities

Types of Attack Vectors 08

Top 3 Attack Vectors 09

Quantity of Attack Vectors 10

Attack Durations 12

Attack Size Distribution 13

Global Attack Source Distribution 14

APAC Attack Source Distribution 15

Global Attack Sources by Autonomous System Number (ASN) 16

Conclusions 17

Research & Methodology 18

NEXUSGUARD®

Q3 2018 Threat Report

Key Observations

The maximum attack size increased 139.84% YoY (Year-over-Year) to 118Gbps, but was down 67.13% QoQ (Quarter-on-Quarter). The average size decreased 81.97% YoY and 96.31% QoQ. As threats eased off from last summer's World Cup peak, total attacks decreased 45.25% YoY and 50.92% QoQ, respectively.

A new development: CSP (Communication Service Provider) networks — especially those at the ASN level — were hit by a stealthy, new volumetric attack whereby attackers contaminate legitimate traffic across hundreds of IP prefixes (some 159 ASNs, spanning 527 Class C networks, based on our findings) with small-sized, junk in order to bypass detection. As a consequence, both maximum and average attack sizes decreased measurably YoY.

By attack vector, SSDP Flood attack counts increased most noticeably, growing more than six-fold from the preceding quarter (more than 120% YoY). We believe the unconventional rise in SSDP Amplification is a result of the new attack pattern targeting CSPs. This pattern also caused the average attack size per IP to fall to only 0.972Gbps during Q3.

Total Attacks

vs. Q3 2017	42.25% ▼
vs. Q2 2018	50.92% ▼

Attack Sizes

118Gbps Maximum Attack Size	vs. Q3 2017	139.84% ▲	vs. Q2 2018 67.13% ▼
0.972Gbps Average Attack Size	vs. Q3 2017	81.97% ▼	vs. Q2 2018 96.31% ▼

DDoS Attack Type

	SSDP	UDP	TCP SYN	ICMP	Application	Amplification
vs. Q3 2017	121.68% ▲	26.38% ▼	68.18% ▼	10.16% ▼	78.13% ▼	26.36% ▼
vs. Q2 2018	639.84% ▲	54.86% ▼	86.28% ▼	45.53% ▼	10.57% ▼	1.65% ▲

Quarterly Focus

DDoS Attacks the “Mongol Way”

As DDoS attack tactics evolve, Communication Service Providers[1] (CSP) at the ASN level are facing a new challenge posed by diffused and stealthy volumetric attacks designed to evade detection. The new tactic resembles the way **Mongol** troops executed battles some 700 years ago. Like the Mongols, today’s perpetrators thoroughly study the targeted landscape prior to mounting their attacks.

By conducting advance reconnaissance to covertly collect information attackers can identify mission-critical IP prefixes. Whereas in the past, attackers tended to zero in on a small number of high-traffic IPs to cause congestion. This sophisticated tactic leads us to believe that such intelligence might be coming from insiders with knowledge of those IP prefixes that are most vulnerable to DDoS attacks.

Mongol military tactics enabled the Mongol Empire to conquer nearly all of continental Asia, the Middle East, and parts of eastern Europe during the 13th and 14th centuries. Highly agile and mobile, horse-riding Mongol soldiers were often sent on scouting missions to gather intelligence about routes and search for terrain most suited to their preferred combat tactics.

1 https://www.gartner.com/it-glossary/csp-communications-service-provider

Junk Traffic Hijacking Legitimate Traffic

Like Mongolian warriors that used human shields, today's perpetrators also use subterfuge to distract and disrupt defenses. In Q3 we observed attacks where perpetrators injected small bits and pieces of junk into legitimate traffic as a disguise. Consequently, attack traffic in the space of each IP address was small enough to bypass detection, but big enough to cripple the targeted site or even an entire CSP network once the traffic converged.

Owing to the negligible size of the junk, typical security devices deployed by ASN-level CSPs are unable to detect and mitigate the traffic before it can cause any harm. This is so because detection thresholds are largely based on the volume of traffic heading to destination IPs.

How is the "bit-and-piece" pattern different from traditional network-layer volumetric attacks?

Bit-and-piece exploits the large attack surfaces of ASN-level CSPs, whereas traditional attacks zero in on one or a few IPs that serve mission-critical services such as websites and mail servers and overwhelm the target by sending voluminous amount of junk. Since the traffic spike is significant and the attack obvious, it's relatively easy to detect abnormalities and mitigate traditional volumetric attacks. In most cases, ISPs with load-balancing capabilities will absorb much of the impact — albeit not 100% — of large volumetric attacks by the time they reach their destination.

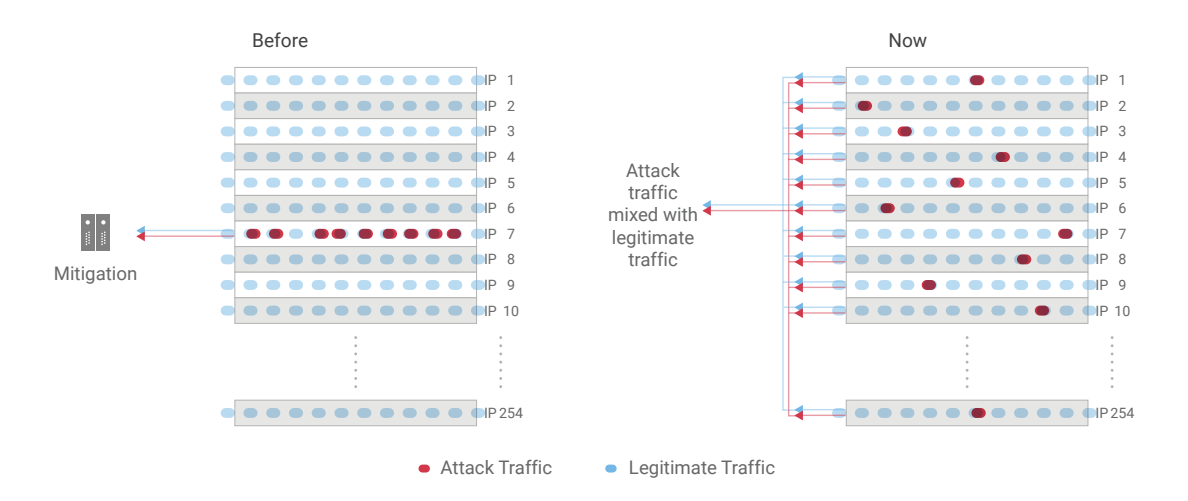

Figure 1. Comparison between Normal Attack Traffic and Attack Traffic with Legitimate Traffic

In the example shown here, the orchestrated attacks generated only 33.2Mbps per destination IP — small enough to fly under the radar and be mistaken as legitimate traffic delivered straight to the destination ASN.

In Q3, Nexusguard observed that some 159 ASNs or 527 Class C networks were targeted in a series of bit-and-piece attacks. The figures reveal that the campaign was significant and the attacks were far more sophisticated than typical network-layer attacks. After tracing advertising BGP routes, AS paths, and trace-routes, we saw that attackers targeted networks within the same geo-location, attempting to max out the physical limitations of transmission lines. In the worst-case scenario outlined in the summary below, the convergence of attack traffic spread across 38 IP prefixes, each loaded with 2.48Gbps of attack traffic — potent enough to overwhelm a 10Gbps ISP line.

Targeted ASNs
159

Total IP Prefixes (Class C Networks) Under Attack
527

Attack Types

- DNS Amplification
- SSDP
- CHARGEN
- NTP Amplification

Targeted Geo-locations

Generally resources physically located within the same geo-location

Ranking	Same Top 10 Attack Campaigns	No. of IP Prefixes in the Same ASNs
1	ISP/Telecommunication	38
2	ISP/Telecommunication	38
3	ISP/Telecommunication	38
4	Datacenter and IP Transit	28
5	Datacenter and IP Transit	26
6	Datacenter and IP Transit	24
7	Datacenter and IP Transit	21
8	Datacenter and IP Transit	21
9	Datacenter and IP Transit	19
10	Datacenter and IP Transit	19

Category	Maximum	Minimum	Average
No. of Targeted IP Addresses per IP Prefix	252	49	131
Attack Durations	1,439.67 min.	5.12 min.	113.81 min.
Attack Sizes per IP	300.1Mbps	2.5Mbps	33.2Mbps
Attack Sizes per IP Prefix	5.32Gbps	285.4Mbps	2.48Gbps

Table 1. Information about Attack Traffic with “Bit and Piece” Pattern

Implications for ASN-level CSPs

Given the negligible size of malicious traffic, targeted ASN-level CSPs can easily miss large-scale DDoS attacks in the making. The diffused traffic is likely to be mistaken as legitimate and delivered straight to the destination ASN. Eventually the ASN will realize its high-traffic IP prefixes are under a multi-gigabyte DDoS attack that is significantly impacting its physical transmission lines.

Black-holing may be a solution. But black-holing all traffic to an entire IP prefix, especially a high-traffic one, will affect large portions of legitimate traffic as black-holing doesn't distinguish between legitimate and malicious traffic. All packets destined for black-holed IP prefixes are dropped, thus disconnecting its upstream networks. And while upstream "clean pipes" may filter many noticeable attacks, the bit-and-piece pattern typically goes unnoticed by upstream ISPs before they converge at the target CSP. In the end, the cumulative impact of junk traffic from diverse IPs creates a severe bottleneck for DDoS mitigation appliances of the targeted CSP. To break the bottleneck, the destination ASN must share the load in order to minimize the impact — for example by multi-casting with a scrubbing facility.

The best solution to mitigating ever-evolving DDoS attacks is an always-on cloud deployment. Nexusguard's global scrubbing centers can be deployed as an always-on solution to mitigate attacks of any size or pattern on the network edge before they reach the CSP.

Attack Vectors and Their Targets

Attack vectors discovered via Nexusguard's honeypot network show that amplification attacks were dominant in the quarter. Simple Service Discovery Protocol (SSDP) Amplification attacks were the most frequent, accounting for 94.1%. CHARGEN came in a distant second at 2.4%, while DNS Amplification followed with a scant 1.8% of attacks observed in Q3.

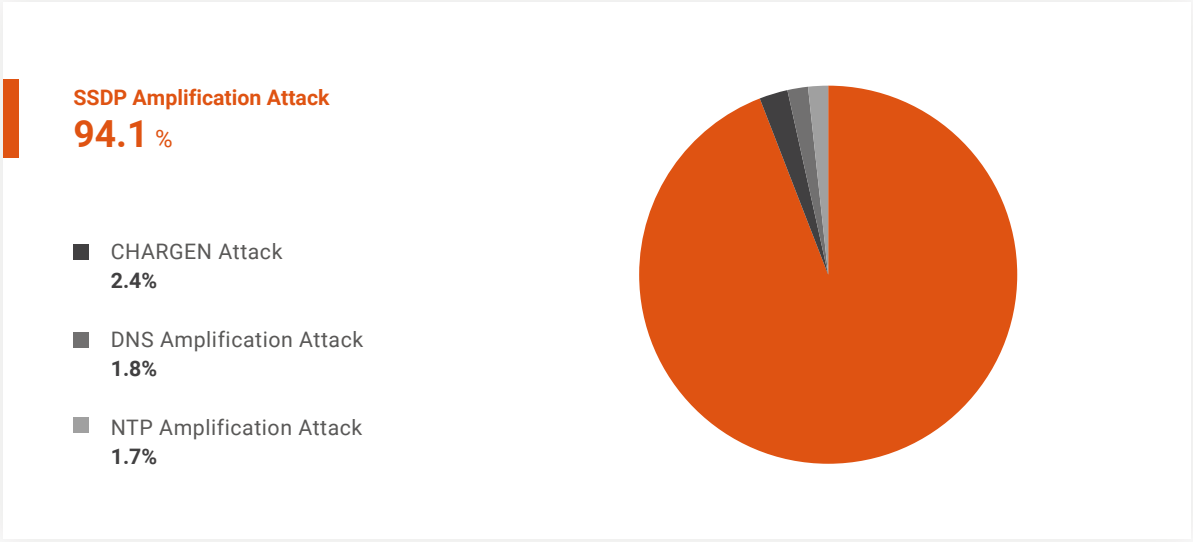

Figure 2. Distribution of DDoS Attack Vectors

Source IP addresses show that ASN-level CSPs were the most popular target in the quarter, accounting for 65.5% of all attacks observed. With so many network assets, including those of their tenants, it's no surprise that ASN-level CSPs are increasingly targeted — directly or indirectly — by DDoS attacks.

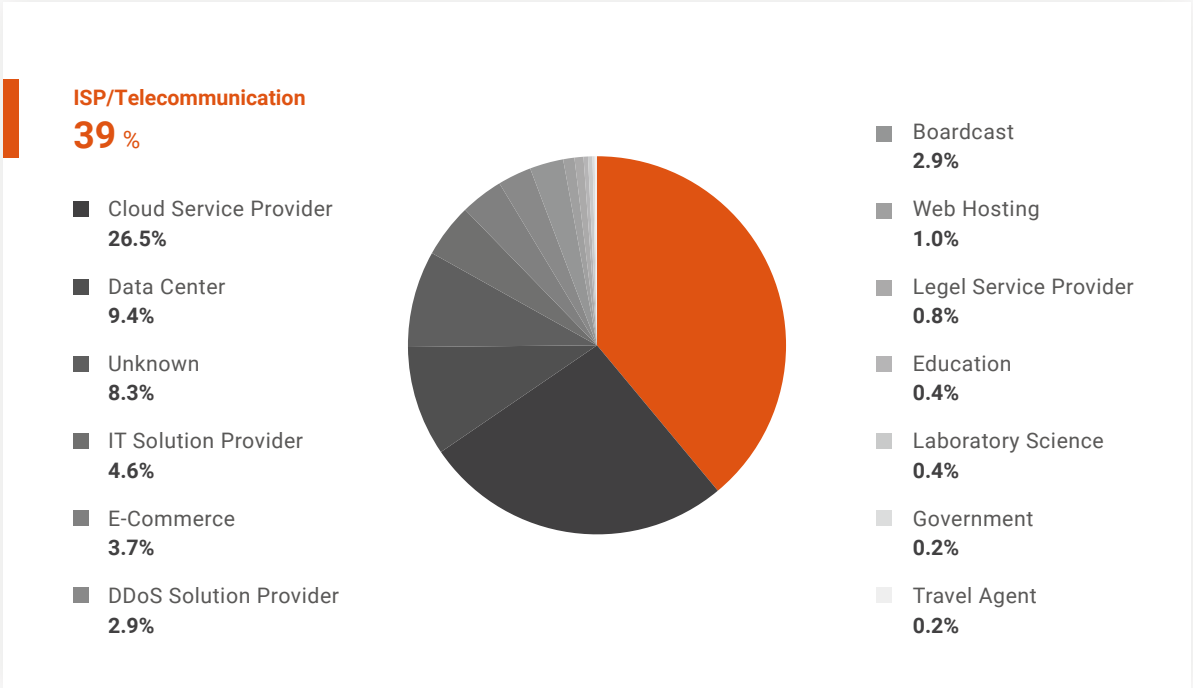

Figure 3. Distribution of Attacks on Different CSP-related Sectors

DDoS Activities

Types of Attack Vectors[2]

SSDP Amplification attacks were the most popular in the quarter, growing 639.84% QoQ and 121.68% YoY, despite the fact that total attack counts fell measurably over both periods. In sharp contrast, UDP attacks fell by 54.86% QoQ and 26.38% YoY and ICMB fell 45.53% QoQ and 10.16% YoY. SSDP attacks totaled 1,820 counts, UDP (1,538) ranked second, while the third and fourth spots were occupied by ICMP (548) and TCP SYN (274).

Because it is open and often unsecured, SSDP is an attractive and vulnerable target. So it's no surprise that attackers abused the protocol to launch "bit-and-piece" DDoS attacks on some 527 Class C networks of CSPs. While SSDP Amplification attacks were the most frequently used in the quarter, Nexusguard believes that attackers will diversify attack vectors going forward.

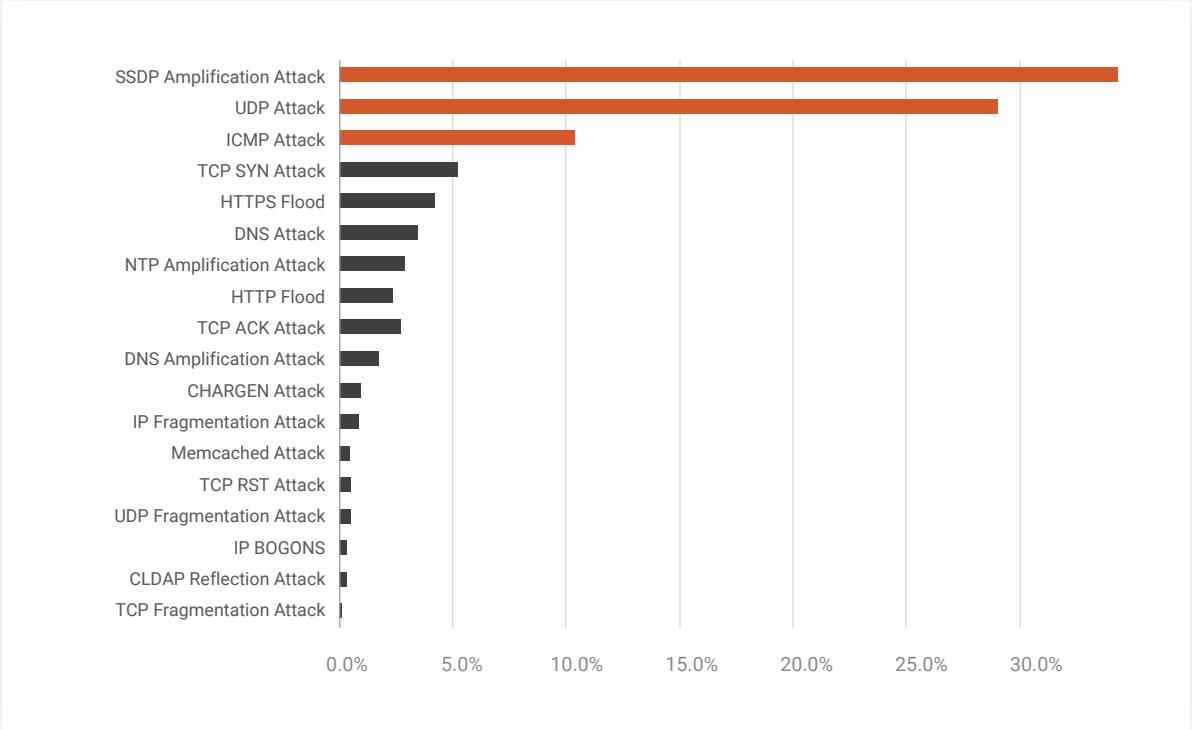

Figure 4. Distribution of DDoS Attack Vectors

2 Attacks on network Layers 3 and 4 lasting for at least five minutes at a size equal to or larger than 100Mbps were counted as volumetric attacks. Attacks targeting applications lasting for at least five minutes with at least 500 requests per sec were counted as application attacks. Attack vector measures the number of vectors exploited by the same attack on the same destination IP. An attack is defined as one attack or more than one attacks that occurred within a time interval of five minutes in between. In the same attack, each attack vector is counted once no matter how many times it is targeted as long as the attacks occurred within a time interval of five minutes in between.

Top 3 Attack Vectors

No.1 SSDP

34.35 %

1,820

SSDP (Simple Service Discovery Protocol) attacks are launched over UDP via Universal Plug and Play devices such as printers, web cameras, routers, and servers. Perpetrators first discover and scan all exploitable devices and then use botnets to send UDP packets with a target's spoofed IP address to UDP Port 1900 of all exploitable devices. In turn, the devices respond massively, causing the target to become inundated with a large volume of replies. According to US-Cert, the bandwidth amplification factor during such attacks can be as high as 30.8x.

No.2 UDP Attacks

29.02 %

1,538

UDP (User Datagram Protocol) attacks can quickly overwhelm the defenses of unsuspecting targets. Speed in detection and response is key to thwarting attackers using this volumetric strategy. UDP frequently serves as a smokescreen to mask other malicious activities such as efforts to compromise personal identifiable information (PII) or the execution of malware or remote codes. When large numbers of UDP packets hit a targeted network, bandwidth is congested and a server's resources sapped, ultimately making them inaccessible.

No.3 ICMP Attacks

10.34 %

548

ICMP (Internet Control Mechanism Protocol) is a connectionless protocol used to diagnose and report errors in networked components such as routers, switches, and IoT devices. Perpetrators send ICMP packets with spoofed IPs to networks with an IP broadcast address. Every device on the network responds with a packet to the victim's IP, causing it to be overwhelmed and preventing legitimate traffic from being handled as it should.

Quantity of Attack Vectors

Nexusguard defines an incident[3] as a series of malicious traffic flows with varying degrees of intensity, regardless of the attack method or signature. A collective analysis of incidents rather than focusing on individual attacks allows us to see the big picture and identify new signatures.

The new "bit-and-piece" attacks we saw in Q3 were mainly abuses of the UDP Port used to generate small-sized SSDP attacks and spread them over a large number of IPs. This stealthy technique is designed to evade detection. We believe attackers will diversify into more attack vectors as they continue to vary this new pattern.

A breakdown of attack vectors revealed that 78.31% of incidents targeted one vector, compared with 52.03% in the previous quarter. 21.69% targeted two vectors or more. Of all multi-vectors analyzed, those that targeted two vectors accounted for 13.44%, while those targeting three accounted for 5.39%. The most complex multi-vector attack targeted as many as ten vectors in a campaign.

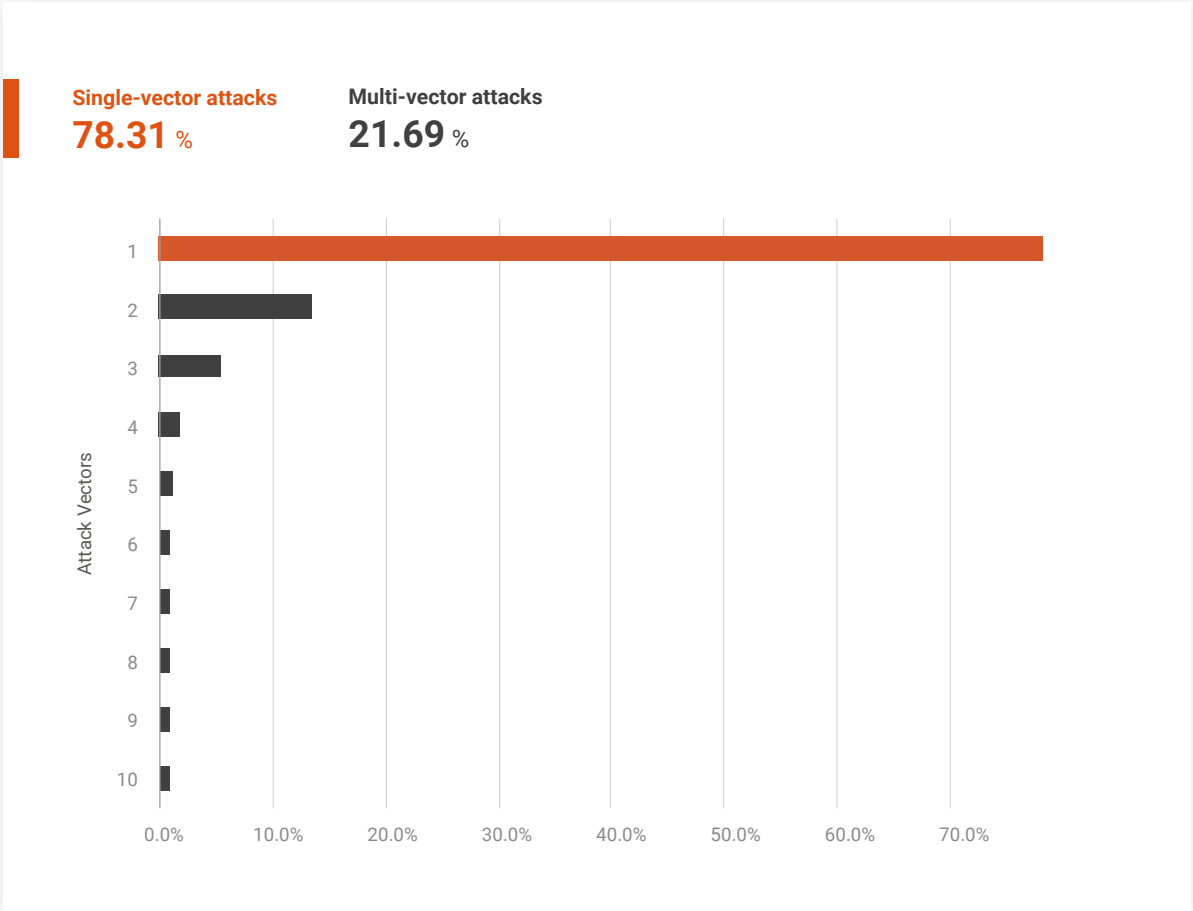

Figure 5. Distribution of DDoS Attack Vectors

3 If more than one attack on the same destination IP are captured and if the time interval between the first and the second attacks is less than 24 hours, both of them will be counted as the same event. If both attacks abuse the same vector, this event will be categorized as a single-vector attack. And if there are more than one attack vector, this event will be categorized as a multi-vector attack.

Multi-vector attacks utilize multiple, simultaneous vectors to maximize the disruption of CSP service availability. UDP was an integral ingredient in each of the top five combinations. The mixture of UDP and DNS was the most popular type of blended attack. The combination of UDP, NTP Amplification, and ICMP ranked number two, while cocktails of UDP, DNS, ICMP, and NTP Amplification attacks were tied for third place.

Rankings	Attack Vector 1	Attack Vector 2	Attack Vector 3	Distribution of Multi-vectors
1	UDP	DNS	N/A	31.45%
2	UDP	NTP Amplification	ICMP	12.26%
3	UDP	DNS	N/A	3.46%
3	UDP	ICMP	N/A	3.46%
3	UDP	NTP Amplification	N/A	3.46%

Table 2. Top Five Multi-vector Attacks

Attack Durations[4]

About 62% of attacks were shorter than 90 minutes, while some 38% lasted longer. Only 0.59% were longer than 1,200 minutes. The average duration was 184.23 minutes, while the longest attack lasted 2 days, 3 hours, and 13 minutes. Shorter, precise attacks enable attackers to maximize disruptions during peak times of online activity in a most cost-effective way.

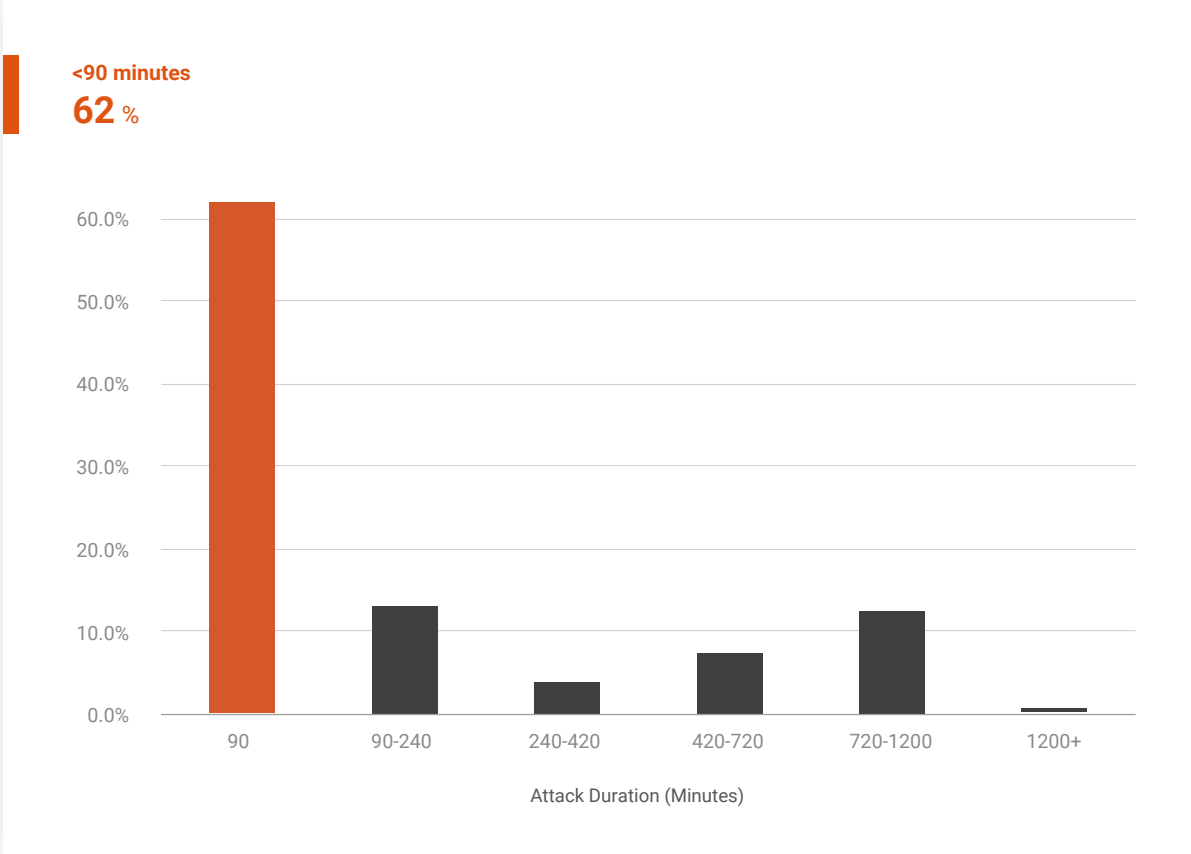

Figure 6. Distribution of Attack Durations

4 Attack duration measures the timespan of a series of attacks on the same destination IP within a time interval of five minutes in between but regardless of the number of attack vectors. If no more attack occurs after five minutes, the finish time of the last attack is considered to be the cut-off time. The "ceasefire breaks" between attacks are excluded from attack duration.

Attack Size Distribution[5]

The average attack size recorded in the quarter was 0.972Gbps. Smaller, new-style attacks (300.1Mbps maximum) were distributed across many IP addresses. Accordingly, most attacks were concentrated in the less than 10Gbps range (91.61%) while those larger than 10Gbps accounted for only 8.39% of the total. While it's true that attacks smaller than 1Gbps are relatively insignificant to large CSP networks, the cumulative impact of bits and pieces of junk traffic distributed across multiple IP prefixes can be substantial when the traffic converges.

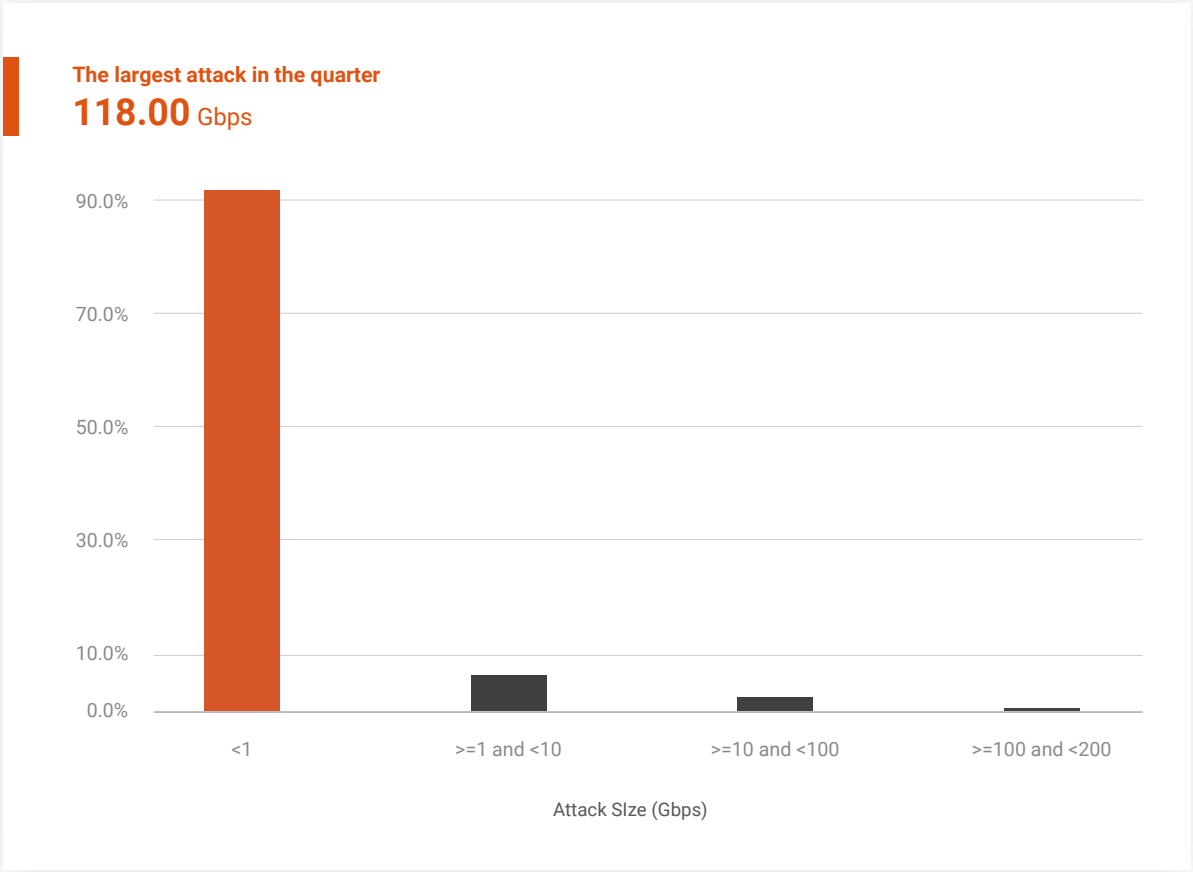

Figure 7. Distribution of Attack Sizes

5 Attack size measures the aggregate size of a series of attacks on the same destination IP within a time interval of five minutes in between but regardless of the number of attack vectors. The peak size of each attack within the same attack is counted in the aggregation. If no more attack occurs after five minutes, the aggregation stops.

Global Attack Source Distribution[6]

China clinched the lead with the largest number of global attack sources. The US followed while Vietnam and Russia placed third and fourth, respectively. China now numbers more than 1B Internet users, nearly one-third of the worldwide total.

Regions	Percentage
China	23.34%
United States of America (US)	14.90%
Vietnam	5.22%
Russian Federation	4.97%
France	4.38%
Brazil	4.34%
South Korea	3.54%
Italy	2.82%
India	2.24%
Egypt	2.18%
Others (135 regions)	32.07%

Table 3. Top 10 Global Attack Sources

6 Untraceable volumetric attacks transmitted with spoofed IP addresses such as TCP SYN, ICMP, and DNS were not included in our sampling. Only traceable attacks like HTTP Flood with real source IP addresses were counted.

APAC Attack Source Distribution

As in the global distribution, China again ranked first. Vietnam followed while India and Thailand took third and fourth place, respectively.

Regions	Percentage
China	59.59%
Vietnam	13.32%
India	5.72%
Thailand	4.57%
Indonesia	3.26%
Taiwan	2.20%
Singapore	2.17%
Hong Kong	1.81%
Japan	1.58%
Malaysia	1.28%
Others (12 regions)	4.50%

Table 4. Top 10 Sources for APAC Attacks

Global Attack Sources by Autonomous System Number (ASN)

The US and China occupied the top three positions with Vietnam and France ranking fourth and fifth, respectively.

ASN	Network Name	Percentage
14061	DIGITALOCEAN-ASN - DigitalOcean, LLC, US	6.21%
45090	CNNIC-TENCENT-NET-AP - SHENZHEN TENCENT COMPUTER SYSTEMS CO LTD, CN	5.54%
4134	CHINANET-BACKBONE - NO.31, JIN-RONG STREET, CN	5.14%
45899	VNPT-AS-VN - VNPT CORP, VN	3.40%
16276	OVH, FR	3.29%
4837	CHINA169-BACKBONE CHINA UNICOM - CHINA169 BACKBONE, CN	2.42%
8452	TE-AS	2.07%
4766	KIXS-AS-KR - KOREA TELECOM, KR	1.85%
42610	NCNET-AS - PJSC ROSTELECOM, RU	1.67%
16509	AMAZON-02 - AMAZON.COM, INC., U.S.	1.49%
Others	1,540 ASNs	66.91%

Table 5. Top Ten ASN Attack Rankings

Conclusions

Owing to their large attack surface, ASN-level CSPs are highly exposed to DDoS attacks. In the third quarter we identified a sneaky, new tactic whereby attackers contaminated a diverse pool of IP addresses across hundreds of IP prefixes (at least 159 ASN, 527 Class C networks) with very small-sized junk traffic. As a consequence, both the maximum and average attack sizes fell measurably YoY.

Like Mongol troops in the past, attackers conducted reconnaissance missions to map out the network landscape in advance and identify the mission-critical IP ranges of targeted CSPs. They then injected bits and pieces of junk into legitimate traffic, which easily bypassed detection because its size was well below detection thresholds.

As opposed to mitigating traffic to a small number of targeted IPs (the traditional volumetric attack method), mitigating broadly distributed, small-sized attack traffic is difficult at the CSP level. The convergence of polluted traffic that slips through the “clean pipes” of upstream ISPs forms a massive traffic flow that easily exceeds the capacity of mitigation devices, leading to high latency at best, or deadlock at worst. Black-holing all traffic to an entire IP prefix may be a way out, yet it is a costly one since black-holing will also block access to a wide range of legitimate services.

The “bit-and-piece” attacks we observed in the quarter often leveraged open DNS resolvers to launch what is commonly known as DNS Amplification, whereby a destination IP (victim) receives only a small number of responses in each well-organized campaign, leaving little or no trace. As such, we expect that it will continue to be difficult to detect and mitigate DNS Amplification attacks carried out in this manner.

Finally, the ongoing evolution of DDoS methods suggests that CSPs need to enhance their network security posture and find better ways to protect their critical infrastructure and their tenants. The continued discovery of new attack patterns should also alert enterprises to the importance of selecting DDoS-proof service providers.

Research & Methodology

As a global leader in Distributed Denial of Service (DDoS) attack mitigation, Nexusguard observes and collects real-time data on threats facing service provider and enterprise networks worldwide. Threat intelligence is gathered via attack data, research, publically available information, Honeypots, ISPs, and logs recording traffic between attackers and their targets. The analysis conducted by our research team identifies vulnerabilities and measures attack trends worldwide to provide a comprehensive view of DDoS threats.

Attacks and hacking activities have a major impact on cybersecurity. Because of the comprehensive, global nature of our data sets and observations, Nexusguard is able to evaluate DDoS events in a manner that is not biased by any single set of customers or industries. Many zero-day threats are first seen on our global research network. These threats, among others, are summarized in quarterly Threat Reports produced by Nexusguard's research team:

- Tony Miu, Research Direction & Security Data Analysis
- Ricky Yeung, Data Mining & Analysis
- Dominic Li, Data Analysis & Content Development
- Jimmy Chow, Technical Writing

20181218-EN-A4

About Nexusguard

Founded in 2008, Nexusguard is a leading cloud-based distributed denial of service (DDoS) security solution provider fighting malicious internet attacks. Nexusguard ensures uninterrupted internet service, visibility, optimization and performance. Nexusguard is focused on developing and providing the best cybersecurity solution for every client across a range of industries with specific business and technical requirements. Nexusguard also enables communication service providers to deliver DDoS protection solution as a service. Nexusguard delivers on its promise to provide you with peace of mind by countering threats and ensuring maximum uptime.

NEXUSGUARD®

www.nexusguard.com