NETSCOUT

ISSUE 8: FINDINGS FROM 2ND HALF 2021

NETSCOUT THREAT INTELLIGENCE REPORT

NETSCOUT Omnis Threat Horizon

If you are looking for real-time data on global DDoS attacks, Omnis Threat Horizon is an invaluable (and free) tool.

NETSCOUT Active Threat Level Analysis System (ATLAS)

NETSCOUT's system for collecting and analyzing internet-wide security events, ATLAS delivers unparalleled visibility into the backbone networks at the core of the internet.

NETSCOUT Atlas Security Engineering Research Team (ASERT)

NETSCOUT's expert researchers and engineers charged with building the tools and front-line database to analyze malware at internet scale, ASERT provides context to the overall threat environment seen within ATLAS.

VIEW LIVE CHART

View any chart in this report as an interactive visualization on our website by clicking the associated button.

Contents

Executive Summary

01

We've got good news and bad news.

The good news is that we accurately predicted a decline in distributed denial-of-service (DDoS) attacks would occur during the second half of 2021 (2H 2021), based on the early decreases we witnessed in Q2 from our last report. These numbers reflect the anticipated decline coinciding with the pre-Omicron easing of COVID-19 pandemic restrictions, with people returning to physical offices and classrooms. The overall number of attacks indeed decreased from 5.4 million in the 1H 2021 Threat Intelligence report to 4.4 million during the second half of the year.

And although any reduction in threat actor behavior is good, the bad news is that the reduction relates only to attacker behavior during the pandemic. Put another way, when you consider attacker behavior independent of the pandemic, the combined total of 9.7 million attacks in 2021 is a 14 percent increase over the number of attacks that occurred in 2019 and represents a DDoS attack every three seconds.

Another good news/bad news scenario also emerged in 2H 2021. The good news was a 32 percent decrease in domain name system (DNS) amplification and a 64 percent decrease in Connectionless Lightweight Directory Access Protocol (CLDAP) amplification attacks, both of which largely account for the overall decrease in attacks for the second half of the year. The bad news is that these types of attacks are now well understood, providing ample incentive for attackers to develop new strategies for disrupting networks and gathering information to extort their targets.

The result is that attackers doubled down on direct-path (non-spoofed) attacks instead of reflection/amplification attacks, evening the playing field between both methods of attack. Likewise, they focused attention on targets that haven't traditionally been in the crosshairs, such as Voice over Internet Protocol (VoIP) providers (who reported an estimated $9 to $12 million in revenue loss), software publishers, and computer manufacturing.

Attackers also started launching more potent direct-path attacks to take down user applications and services, thereby disrupting consumers' ability to access the internet. Meanwhile, they continued to innovate with server-class botnets and increased use of DDoS techniques such as carpet-bombing.

So although it's tempting to simply look at the decrease in overall attacks as threat actors resting on their laurels, the reality is that attackers are innovating and adapting new techniques and methodologies to strengthen and monetize their nefarious behavior.

Key Findings

01

The Triple Threat

For the first time ever, three prolific DDoS extortion campaigns operated simultaneously. VoIP providers were pummeled with high-profile DDoS extortion or ransom DDoS (RDDoS) attacks from a REvil copycat, resulting in an estimated revenue loss of $9 to $12 million, while Lazarus Bear Armada (LBA) and Fancy Lazarus targeted organizations around the world.

Meanwhile, ransomware gangs continued adding triple extortion—attacks made up of file encryption, data theft/leakage, and DDoS attacks—to their arsenals.

02

A Flood of Attacks

Adversaries inundated organizations with TCP- and UDP-based floods, an activity we refer to as direct-path (non-spoofed) attacks. Nevertheless, a decrease in some amplification attacks drove down the total attack count for 2H 2021.

TCP-BASED FLOODS SURPASSED SOME REFLECTION/ AMPLIFICATION

-32% Decrease in DNS amplification attacks, resulting in an overall 14% decrease in attacks from 1H 2021

-64% Decrease in CLDAP amplification attacks

03

DDoS Ripple Effect

A rise in industry-specific targeting and direct-path attacks indicates that adversaries ramped up targeting of organizations, while attacks targeting customers of internet service providers (ISPs) on wired and cloud hosted networks declined slightly. This shift in modus operandi largely began as the world resumed normal daily activities in August and September 2021, coinciding with schools resuming on-site classes, companies removing some COVID restrictions, and employees returning to the office. Despite these focused targets, DDoS attacks cause damage not only to the intended target but to everything around it.

ATTACKERS ZEROED IN ON A NUMBER OF INDUSTRIES FOR DDoS ATTACKS, INCLUDING:

- **VoIP** 93% increase
- **Electronic Computer Manufacturing** 162% increase
- **Computer Storage Device Manufacturing** 263% increase
- **Software Publishers** 606% increase
- **Insurance Agencies and Brokerages** 257% increase
- **Colleges, Universities and Professional Schools** 102% increase

04

The Rise of Server-Class Botnet Armies

The first botnets in the early 1990s were composed of servers, followed over the years by general-purpose personal computers (PCs) and then Internet of Things (IoT) botnets, which rose to prominence in the 2010s. Recently, adversaries not only increased the size of IoT botnets but also conscripted high-powered servers into larger botnets, as seen with the GitMirai variant exploiting a vulnerability on Git Servers.

05

DDoS-for-Hire Free-for-All

Launching DDoS attacks with illicit DDoS-for-hire services no longer requires even a nominal fee. Most services now allow users to test basic DDoS attacks before increasing attack potency via some form of digital or cryptocurrency. The range of services offered by these nefarious platforms spans layers 3, 4, and 7 and targets everything from specific applications and games to methods for bypassing standard anti-DDoS measures. According to just 19 out of hundreds of such sites on the dark web, they claim to have successfully launched more than 10 million DDoS attacks.

06

The Intersection of Encryption, State, and DDoS Defense

Adversaries are laser-focused on disrupting layer 4 Transport Layer Security (TLS)-encrypted applications and services, evidenced by the increase in bandwidth and throughput of these attacks. In fact, DDoS-for-hire services increasingly added specific attack types for different web browsers, web-based games, and gaming services software. These attacks negatively impact stateful firewalls, load balancers, and intrusion prevention systems (IPSs), further emphasizing that DDoS attacks are attacks against capacity or state.

Attack History

1H 2021 **5,351,930**

\+

2H 2021 **4,406,713**

→ **9,758,643**

3% decrease over 2020 / 14% increase over 2019

02

Global DDoS Attack Trends

The second half of the year brought about the establishment of high-powered botnet armies and a rebalancing of the scales between volumetric and direct-path attacks, creating new standard operating procedures (SOPs) for attackers and adding new tactics, techniques, and procedures (TTPs) to their arsenals.

This was observed as TCP-based flood attacks like TCP SYN, ACK, and RST floods remained stable, while DNS and CLDAP amplification attacks decreased by 32 percent and 64 percent respectively. The decrease in DNS and CLDAP amplification resulted in a return to prepandemic attack counts for 2H 2021 at 4,406,713 attacks. This represents a 14 percent decrease from 1H 2021 but a two percent increase from 2H 2019.

Global Stats: Number of Attacks

4,406,713

14% decrease from 2H 2020

Average Attack Duration
51 minutes (31% increase)

Largest Attack

612 Gbps *

14% increase from 2H 2020

Date
November 6, 2021

Target
Czechia

Vectors Used
DNS, DNS amplification, ICMP, TCP ACK, TCP RST, TCP SYN

Attack Duration
16.83 minutes

Fastest Attack

453 Mpps

107% increase from 2H 2020

Date
December 7, 2021

Target
Russia

Vectors Used
CLDAP amplification, ICMP, TCP ACK, TCP RST, TCP SYN, TCP SYN/ACK amplification

Attack Duration
1 hour 44 minutes

*

NETSCOUT observed multiple terabit-class attacks during the year; however, we only list maximum attacks for which we observed the totality of the attack rather than partial data.

Despite the observed decrease in amplification attacks in 2H 2021, the year ended with 9.7 million DDoS attacks in total (an attack every three seconds!), a mere 3 percent decrease from the record number of attacks that took place during the height of the pandemic. This clearly signals that it would be premature to roll the victory drums, given the clear and present dangers lurking in the DDoS threat landscape.

This ebb and flow in DNS amplification attacks is a trend that tracks back to 2018. Similar dips occurred in May 2018, February 2019, September 2019, July 2020, and June 2021. Despite the occasional drop in overall attack numbers, the trend maintained an up-and-to-the-right trajectory at the close of December 2021.

A month-to-month comparison from 2H 2019 to 2021 illustrates how the pandemic impacted DDoS activity, including peaks occurring in January and March 2021. A decrease in attacks against consumers on wireline ISP networks sharply contrasts with a marked increase in attacks against education, computer and software manufacturing entities, and wireless telecommunications providers. This likely is due to multiple factors, including a return to in-person education and the rapid adoption of 5G wireless technology.

9.7M

Attacks in 2021, a mere 3% decrease from the record number of attacks that took place during the height of the pandemic

Monthly DDoS Attack Frequency 2H 2019 to 2H 2021

VIEW LIVE CHART

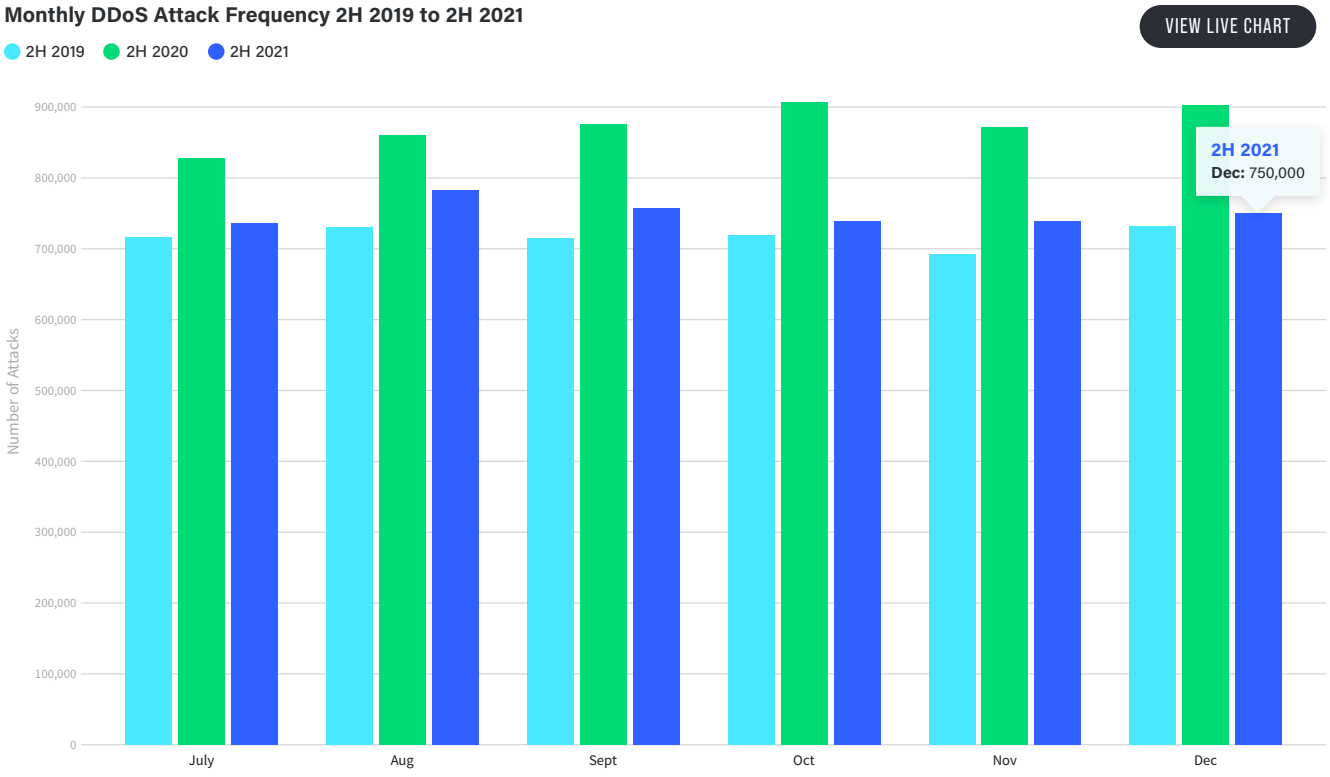

Figure 1: Monthly DDoS Attack Frequency 2H 2019 to 2H 2021 (Data: Omnis Threat Horizon)

DDoS Extortion and The Triple Threat

7-Figure Losses from DDoS Attacks Reported by Publicly Traded Company

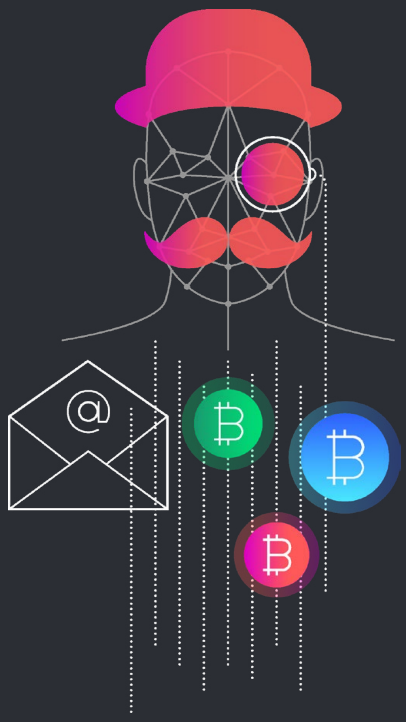

Although DDoS extortion (aka RDDoS) isn't new, high-profile DDoS extortion attack campaigns sometimes emerge. It's not unusual to have one high-profile DDoS extortion campaign in a year, but it's fairly rare to see two such campaigns in a year. During 2021, however, a new record was established as three high-profile DDoS attack campaigns took place. This also signals that ransomware gangs are laser-focused on increasing the use of triple-extortion attacks (ransomware + data theft + DDoS).

The prolific Lazarus Bear Armada (LBA) DDoS extortionist threat actor extended its high-impact attack campaign into 2021, targeting multiple verticals worldwide and exhibiting a high degree of pre-attack reconnaissance to maximize attack efficacy.

The Fancy Lazarus DDoS extortionist kicked off a campaign that initially targeted the authoritative DNS servers of wireline broadband access ISPs in the U.K. and Scandinavia by using DNS reflection/amplification attacks, a suboptimal vector when attacking authoritative DNS servers. The campaign was somewhat successful due largely to the unpreparedness of a few network operators; nevertheless, the attacks were mitigated relatively quickly.

The third high-profile DDoS extortion campaign of the year was an aggressive series of attacks masquerading as the REvil ransomware group and targeting SIP/RTP VoIP operators. Retail and wholesale VoIP providers in the U.K. were the initial targets, followed by attacks against VoIP operators in Western Europe and North America. Notably, one VoIP wholesaler filed a form with the U.S. Securities and Exchange Commission (SEC) estimating the total cost of the DDoS attack at between $9 and $12 million. Attackers now appear to view DDoS attacks as criminal endeavors in and of themselves—as opposed to one pillar of triple extortion attacks—meaning more-skilled DDoS extortion campaigns should be expected as sophisticated ransomware groups master this tactic.

Ransomware Gangs

In the 1H 2021 Threat Intelligence report, we noted that several different groups conducting ransomware operations have also moved into DDoS attack territory to place greater pressure on victims to pay demanded ransoms. For this report, Palo Alto's Unit 42, a Threat Intelligence partner, created a summary of active and recently inactive ransomware gangs that also use DDoS to extort victims into paying the ransom. The following groups are known to use and have been observed using DDoS as part of their operations.

paloalto NETWORKS

UNIT 42 BY PALO ALTO NETWORKS

A NETSCOUT PARTNER

Unit 42 is a premier threat intelligence and cybersecurity consulting organization chartered to identify and resolve the most challenging threats and make the world a safer place.

Avaddon

Avaddon ransomware was first seen in February 2020 and by June 2020 had quickly evolved into ransomware as a service (RaaS). In January 2021, the group evolved again to include DDoS attacks in its extortion repertoire. Despite a successful run, the group inexplicably shut down its operation in June 2021, possibly as a result of political pressure and/or the release of private keys that enabled victims to decrypt files.

REvil

Although currently not operational due to a global takedown, REvil was a prominent user of RaaS. With its highly adaptable encryptors and decryptors, REvil provided infrastructure and services for communicating with victims, as well as a leak site for releasing stolen data if the victim refused to pay the ransom. In February 2021, REvil announced that it would begin contacting its victims' business partners and the media to disclose breaches and further extort victims. On March 5, 2021, a REvil spokesperson announced the addition of DDoS attacks, effectively elevating the group's TTPs to include multi-extortion.

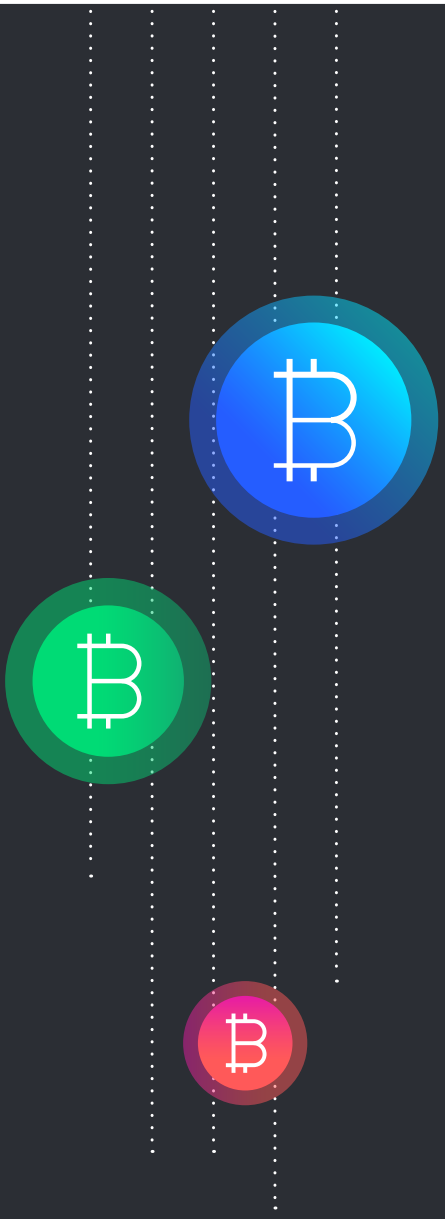

BlackCat

One of the newest ransomware groups, BlackCat (aka ALPHV), was discovered in November 2021. Operating as a RaaS, the group quickly gained notoriety for its sophistication and innovation. BlackCat solicits for affiliates in known cybercrime forums by promising to leverage ransomware and give 80 to 90 percent of the ransom payment to the affiliate, with the remainder paid to the BlackCat author. The malware itself is written in Russian and coded in Rust, making it one of the first pieces of ransomware to use it. BlackCat not only encrypts and steals victims' data, but it also then threatens to leak the data via a leak site. Should the victim need additional persuasion to comply with the ransom demand, BlackCat threatens a DDoS attack.

AvosLocker

First seen in summer 2021, AvosLocker is simple but effective ransomware that has utilized triple extortion from the start. AvosLocker operators advertise in underground networks for affiliates with active directory experience, as well as for "access brokers" who potentially could provide access to compromised systems. Affiliates are incentivized with having AvosLocker take care of the extortion and negotiation parts of the process. AvosLocker then uses affiliates to infect a victim, while handling the remaining ransomware process itself. Like some other ransomware groups, AvosLocker operates a leak site to apply additional pressure on victims to pay the ransom. The group has attacked a diverse set of victims in terms of both region and industry.

Suncrypt

Initially appearing in October 2019, Suncrypt was one of the first ransomware groups to launch DDoS attacks. Along with data encryption and theft, Suncrypt extorts its victims by threatening to attack infrastructure or networks. Likewise, further pressure is applied by threatening to expose the breach to employees, stakeholders, and the media should ransom negotiations fail. The group maintains a leak site and promises that it won't expose victim data during the negotiation process. If that process fails, however, Suncrypt leaks victim data and initiates a DDoS attack until negotiations resume.

Botnet Army Adds New Weapons

The commonly held idea of botnets used for launching DDoS attacks is that compromised IoT devices come under the control of attackers via a common command-and-control (C2) infrastructure.

The first DDoS-capable botnets debuted in 2007, and they became commonplace by 2013. Their popularity soared in 2016 after the source code of the Mirai IoT botnet was leaked. In 2H 2021, they continued evolving with the convergence of Mirai for Intel architectures, which inadvertently resulted in the rapid exploitation of serious vulnerabilities in servers running Confluence, GitLab, and Log4J. Exploits were crafted and delivered to compromise significant numbers of powerful, highly connected servers that were brought together via standard botnet C2 architectures.

Given that online criminals are familiar with the DDoS capabilities of existing Mirai botnets, they were able to quickly employ the new server-class Mirai botnets to launch vicious DDoS attacks. In 2H 2021, two direct-path flooding attacks of more than 2.5 Tbps were launched using server-class botnets. These are the first known terabit-class, direct-path DDoS attacks; previously, reflection/amplification attacks were considered the most practical way to launch DDoS attacks of this magnitude.

The newfound popularity of server-class DDoS botnets is linked with the growth in direct-path DDoS attacks, when compared with reflection/amplification attacks. We expect this trend to continue, driven by the introduction of multigigabit consumer wireline and wireless 5G broadband internet connectivity, increasingly powerful home computers, and IoT devices. We also foresee the very definition of server-class nodes expanding beyond the internet data center (IDC) and into the residential space.

The Dark Side of DDoS-for-Hire

The dark web is a dangerous place where adversaries own and operate DDoS-for-hire platforms and botnets to launch everything from free tests to high-powered multivector attacks. ASERT explored this underground space to evaluate the kinds of attacks being launched. Likewise, we wanted to better understand the kinds of platforms used and their capabilities, to illustrate the low barrier to entry and why DDoS attacks are so prevalent.

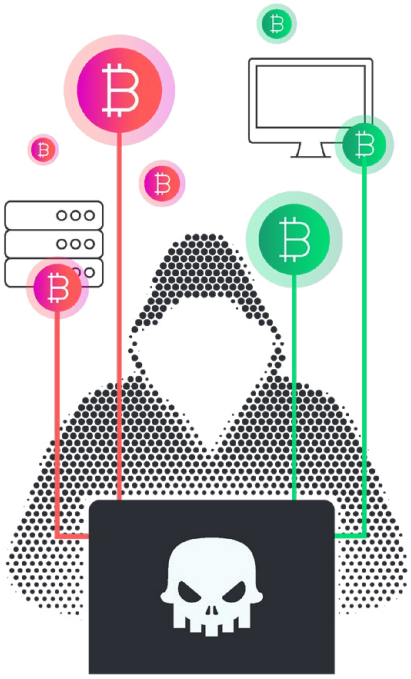

As such, we researched the top 19 validated DDoS-for-hire services and captured the types of attacks, purported number of users, and the costs to launch attacks.

1. AnonBot
2. Booter
3. Booter SX
4. CryptoStresser
5. CyberVM
6. DDoS Service
7. Downed
8. FlyStress
9. Instant Stresser
10. IPStresser
11. NetworkStress
12. Project Delta
13. Str3ssed
14. Stresser GG
15. Stresser US
16. SunStresser
17. Toxicity
18. WebStresser
19. ZDStresser

Although some of these services have static pricing models, many of them allow for custom configurations based on duration, concurrent tests, and power, which is how adversaries measure bandwidth and throughput.

Prices for these services vary wildly. We found free tests, tests for $5 over a five-day trial, and full attacks for as much as $6,500, which included 100 concurrent attacks, no daily limits, and a committed 1 million packets per second (Mpps). NetworkStress service boasts a 1 Tbps attack size using 150,000 bots for $2,499. Although these services boast massive capacity, we have yet to observe any DDoS attacks sourced from them in the terabit range.

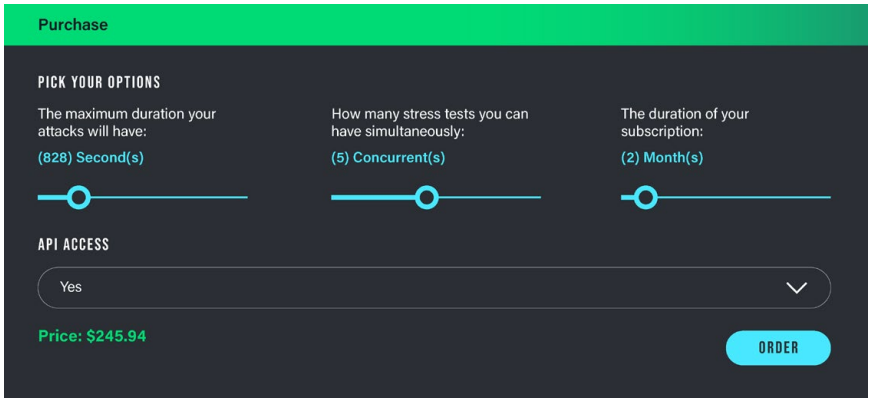

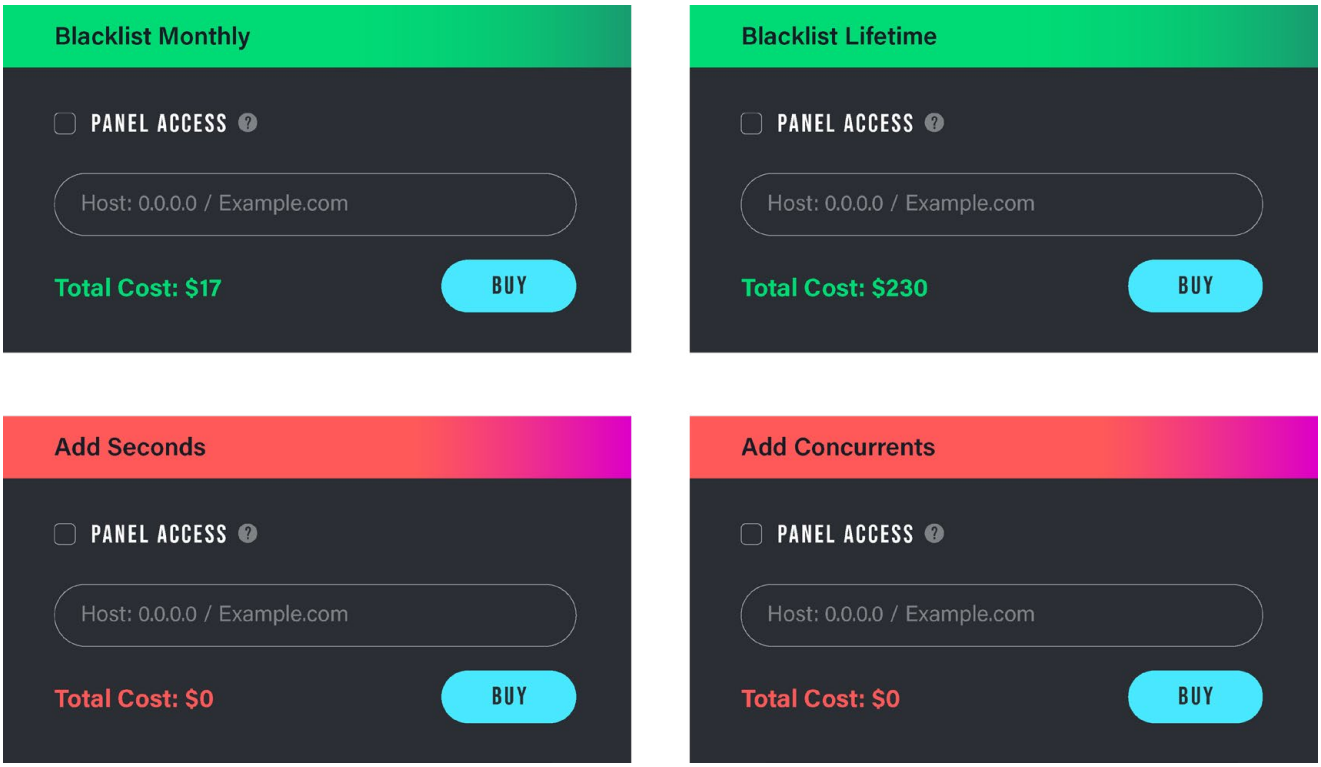

In the 1H 2021 Threat Intelligence report, we described how some of these underground services offer "blacklists" or delisting services to prevent attacks. One example of this can be found on Booter SX, where adversaries offer a temporary or permanent option for delisting IPs. At least three of the services noted above include this feature, which is anything but a guarantee the purchaser will not be attacked.

Nearly every service offers some form of free DDoS attack capability via Network Time Protocol (NTP), DNS, CLDAP, or a random UDP reflection/amplification attack vector. In addition to the free options, these 19 platforms combined boast a total of more than 200 different attack types, many of which are shared across platforms. UDP and TCP reflection/amplification are the most prevalent, followed by UDP and TCP floods. The services also offer varying degrees of UDP and TCP bypasses for CAPTCHAs or other anti-DDoS defenses.

Despite the incredible diversity of these platforms, the majority of attack types are recognized and predominantly mitigated via standard defensive practices. Our primary motivation in exploring these services was to determine the capabilities available to adversaries. Based on our research, none of the listed services was a surprise or provided something we haven't witnessed in the wild. Given a solid understanding of these attack methods and a properly tuned mitigation platform, network security professionals can create defensive measures and templates to counter attacks from booter/stresser services.

ATTACK TYPES OFFERED ON DDOS-FOR-HIRE PLATFORMS

- COAP amplification
- OVHGameTCP
- NTP amplification
- SNMP amplification
- SynAck
- DNS amplification
- CF-Bypass
- LDAP amplification
- WSD amplification
- DVR amplification
- HTTP
- CLDAP amplification
- ESP Flood
- SSDP amplification
- TCP
- FIVEM
- MixUDPAMP
- UDP
- SOURCE
- VSE amplification
- ARM amplification
- MINECRAFT
- IPSec
- HEAD
- GoogleCloud

The Intersection of Encryption, State, and DDoS Defense

Application-Layer DDoS Attacks Versus DDoS Attacks Against Applications

One of the most important and wide-reaching trends in the security landscape over the past decade has been the industrywide push to implement strong encryption for websites, online applications, communications services, and just about everything else we use online.

This wholesale move toward encryption for anything and everything also has been noted by attackers. The additional overhead required to process encrypted communications at large scale often means that launching successful DDoS attacks against encrypted applications and services requires comparatively fewer resources on the part of the attackers. Conversely, DDoS defense for encrypted applications and services also requires more resources on the part of defenders.

High-volume application-layer attacks launched over HTTP/S were prominent during this period. Attacks launched via the Meris and Dvinis router-based botnets were reported, either originating directly from the bots themselves or being relayed through them by way of the SOCKS5 proxy functionality of the bots. Attacks of up to 17.2 million requests per second (Mrps) were reported, representing a significant new metric for HTTP/S-encrypted application-layer DDoS attacks.

Looking at a two-year snapshot for bandwidth and throughput in attacks targeting applications and services on TCP port 443, we see significant trends toward more potent attacks.

17.2 Mrps

Although not an uncommon benchmark in volumetric DDoS attacks, this was one of the fastest throughput attacks observed to date from a botnet directed at HTTP/S-encrypted application-layer services.

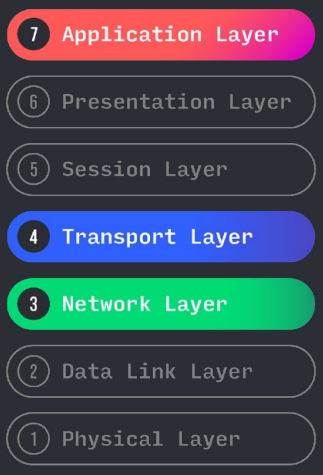

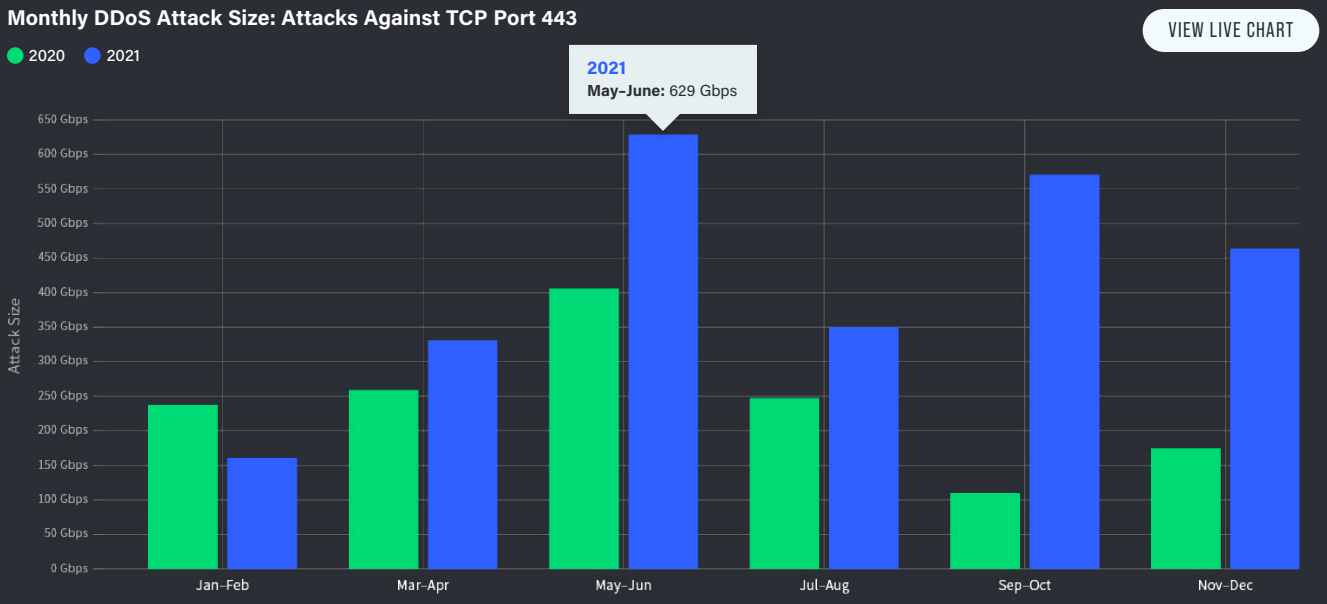

Figure 2a: Monthly DDoS Attack Size (Data: Omnis Threat Horizon)

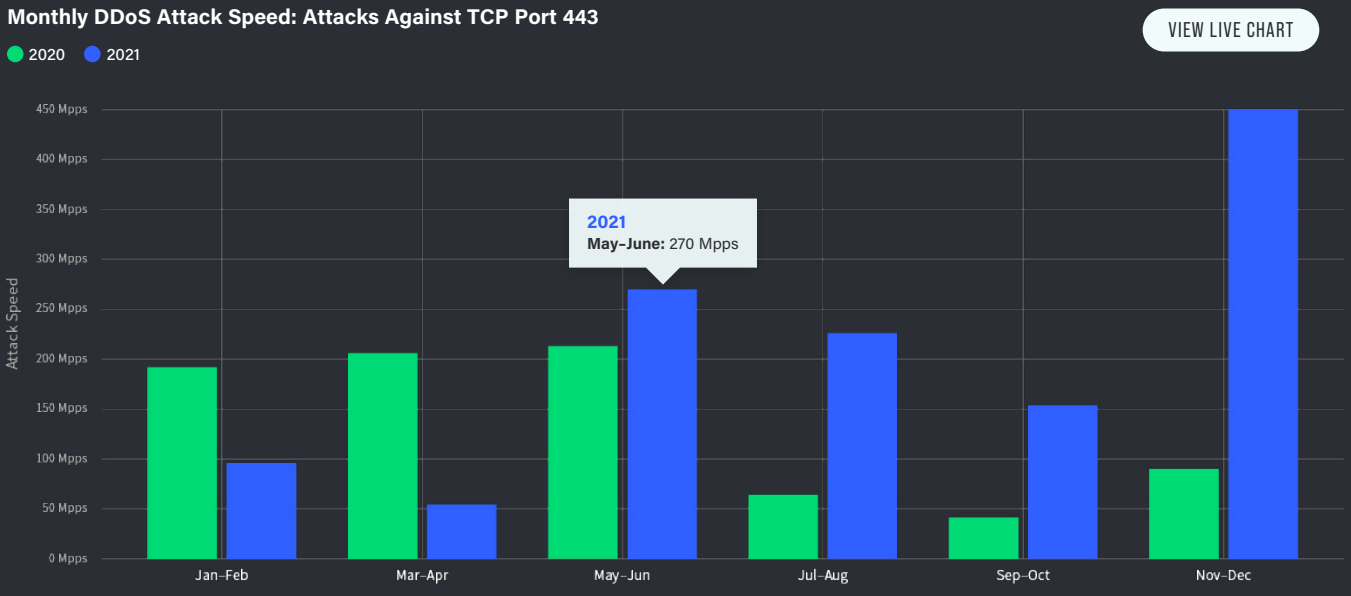

Figure 2b: Monthly DDoS Attack Speed (Data: Omnis Threat Horizon)

It is ironic that measures intended to bolster two aspects of security—confidentiality and integrity—can have unintended consequences for security's third (and arguably most important) aspect: availability. Although it is important that deployment of TLS 1.3 proceeds apace, organizations must take into account the associated increases in complexity and overhead, while ensuring that their public-facing properties are designed and implemented to minimize state and maximize DDoS defense capabilities, thereby ensuring maximal resiliency in the face of attack.

Carpet-Bombing Attacks

Our Threat Intelligence partner Neustar also witnessed a significant shift in an attack methodology, with carpet-bombing picking up steam in July 2021. This attack is akin to flinging sand instead of a rock with the hope that many smaller attacks will succeed where a single, large attack fails. Data from Neustar's security operations center (SOC) revealed that carpet-bombing attacks outnumbered individual attacks by more than 10 percent in 2H 2021. The very nature of these attacks makes them difficult to defend against, because there are multiple points to protect as opposed to a single point of entry.

Such attacks, which often can be too small on their own to trip mitigations, can cause a host of distractions and confusion as they land across a target's network. They certainly make it necessary for defenders to update detection mechanisms and policies to spread defenses across all externally facing ingress points.

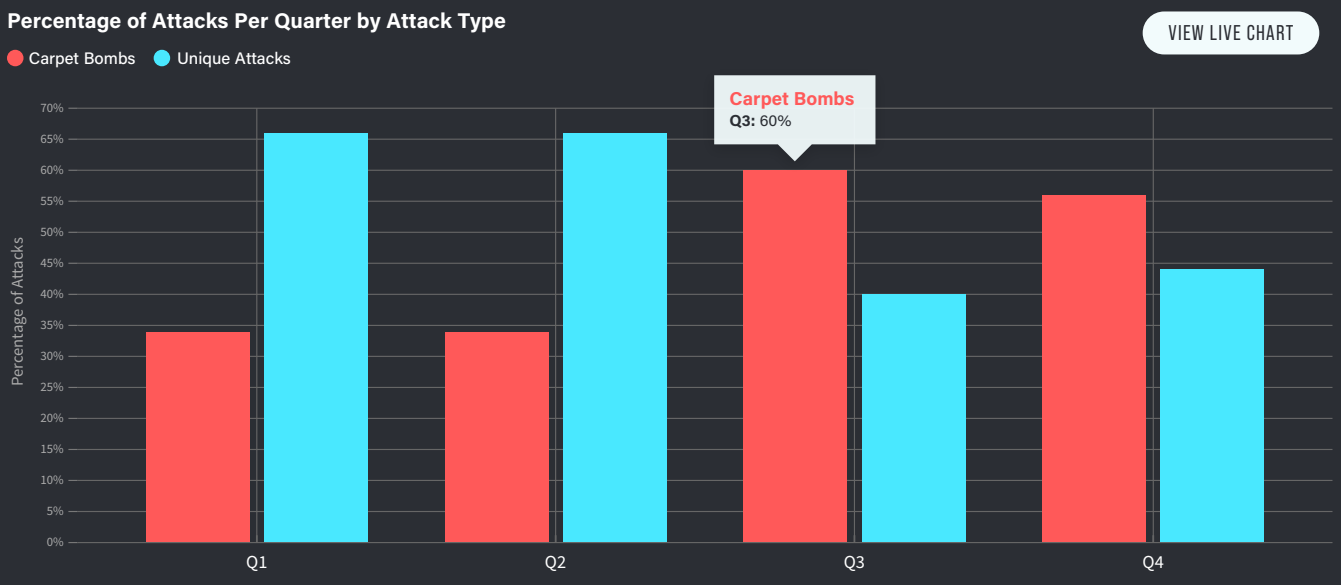

Figure 3a: Percentage of Attacks Per Quarter by Attack Type (Data: Neustar SOC Data)

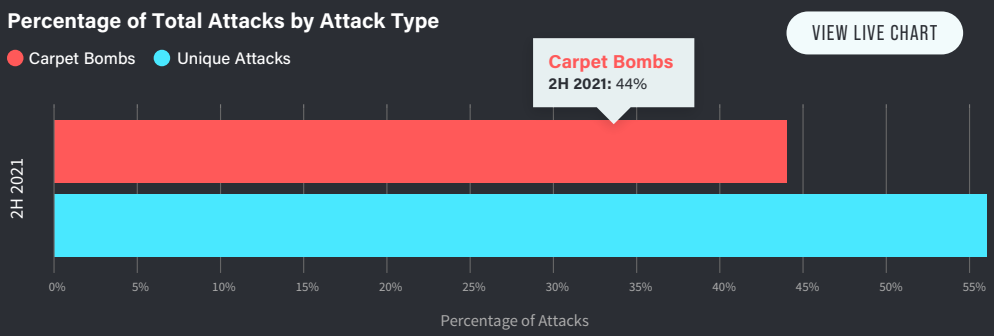

Figure 3b: Percentage of Total Attacks by Attack Type (Data: Neustar SOC Data)

Vertical Industries

Always a popular target for attacks, many of the telecommunications verticals nevertheless saw fewer attacks in 2H 2021. One of the more notable exceptions occurred in the wireless telecommunications space, where a likely increase in wireless hotspot gaming and the rapid adoption of 5G fueled increased attacks (see Industry Spotlight: Wireless Telecommunications). Meanwhile, the closely related software and computer manufacturing verticals witnessed massive increases in attacks (see Industry Spotlight: Digital Supply Chain).

As adversaries sought to cash in on DDoS extortion, they increasingly launched attacks against insurance agencies and brokerages (see Industry Spotlight: Insurance Agencies and Brokerages), as well as against VoIP providers (see Industry Spotlight: VoIP Providers). Unfortunately, some of these attacks were highly successful, causing significant damage both to the targeted organization and collaterally with their customers.

We'd be remiss in not mentioning one more motivation that stands the test of time: "Because I Can." Some people like to watch the world burn. And because they can, they do. By fall 2021, people were returning to normal life, including a return to physical versus virtual classrooms. The increase in attacks on colleges, universities, and professional schools is likely attributable to students looking to start fires wherever possible (see Industry Spotlight: Colleges, Universities, and Professional Schools).

METHODOLOGY

Vertical industry discussion is based on analysis of attack data by North American Industry Classification System (NAICS) codes, which group companies into 22 broad categories that contain multiple large subvertical sectors.

The data represented in our vertical analysis represents less than one quarter of our attack counts and should be viewed as a sampling of our larger dataset.

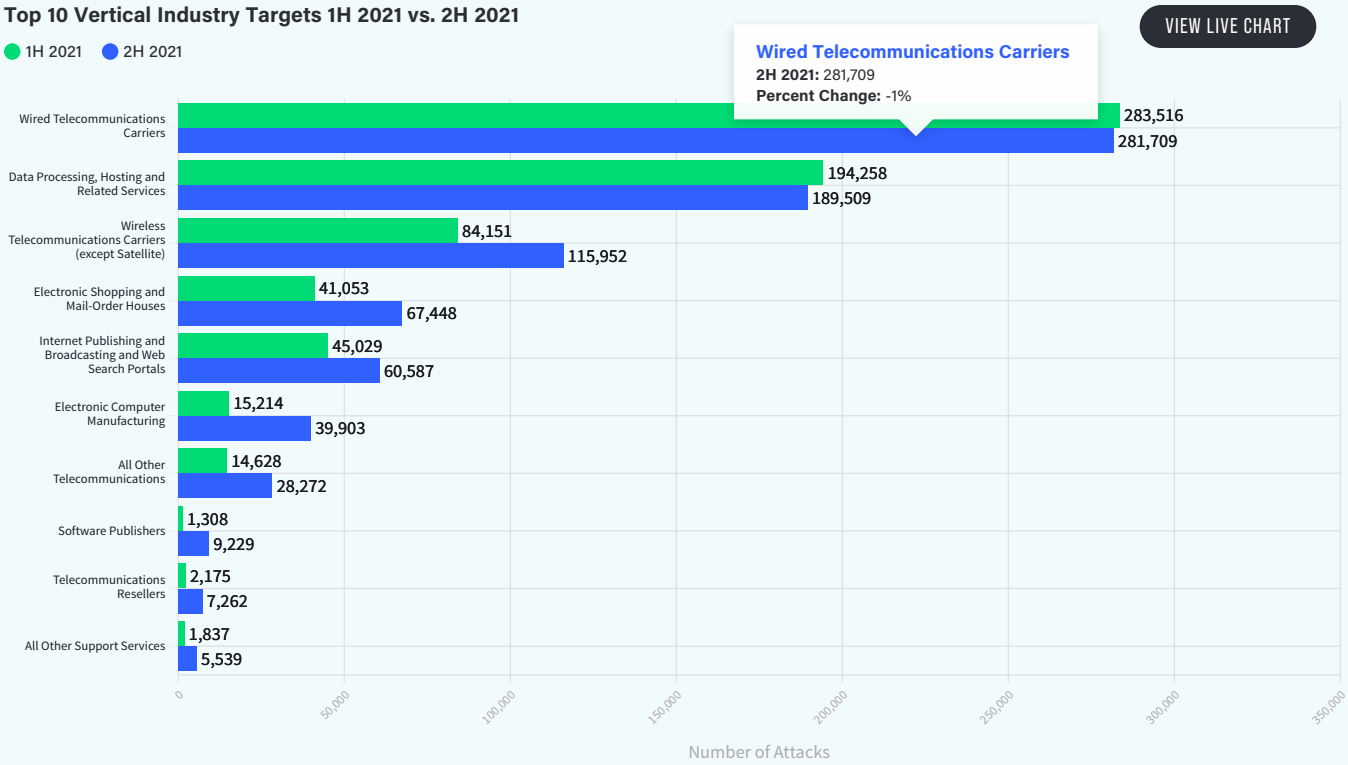

Figure 4: Top 10 Vertical Industry Targets 1H 2021 vs. 2H 2021 (Data: Omnis Threat Horizon)
Note: This is a sampling of our dataset.

Industry Spotlights

WIRELESS TELECOMMUNICATIONS

Gamers received a small breath of fresh air as DDoS attacks against consumers on wireline networks saw a mild decrease. Sadly, this reprieve for one type of consumer shifted to an increase for wireless consumers. The wireless industry experienced a disproportionate increase in attacks—even as many other telecommunications types saw declines during 2H 2021. This trend likely reflects a continued increase in gamers leveraging wireless hotspots and the rapid expansion of 5G technologies and services.

Historically, we've seen a larger share of DDoS attacks against this segment in Asia Pacific (APAC); however, for the second half of the year, we instead observed a 38 percent increase in DDoS attacks globally.

DIGITAL SUPPLY CHAIN (SOFTWARE PUBLISHERS AND COMPUTER MANUFACTURING)

We observed a 606 percent increase in attacks against software publishers compared with 1H 2021. Combined with a 162 percent increase in attacks on computer manufacturers and a 263 percent increase against computer storage manufacturing, it becomes apparent that attackers are focusing a concerted effort on the digital supply chain.

COLLEGES, UNIVERSITIES, AND PROFESSIONAL SCHOOLS

Although DDoS extortion and attacking gamers for monetary gain are the top motivations behind DDoS attacks, we sometimes see attacks that are designed by students who want to play hooky or delay a test. Such was the case in 2H 2021, when attacks against colleges, universities, and professional schools increased by 102 percent. These attacks coincided with a return to physical classrooms, and they serve as a stark reminder to educational institutions that they can easily fall prey to DDoS attacks that can have significant impact on both faculty and the student body.

VoIP PROVIDERS (ALL OTHER TELECOMMUNICATIONS AND CLOUD PROVIDERS)

DDoS extortion campaigns also resulted in numerous VoIP providers all over the world being taken offline. VoIP providers and their infrastructure fall under two primary verticals as defined by the North American Industry Codes: all other telecommunications, and data-processing hosting and related services (cloud computing). The first of these had a 93 percent increase in attacks from 1H 2021, and the second saw a notable increase in Europe, the Middle East, and Africa (EMEA), where most of these attacks occurred. In fact, the data-processing hosting and related-services category was the top target in EMEA for 2H 2021.

INSURANCE AGENCIES AND BROKERAGES (DDoS EXTORTION)

Insurance agencies and brokerages—always a favored target for DDoS extortion attacks—experienced an increase in attacks of 257 percent compared with 1H 2021. This segment was an early target for the LBA campaign dating back to mid 2020.

DDoS Attack Vectors

VIEW LIVE INTERACTIVE PERIODIC TABLE

DNS Amplification

A DNS reflection/amplification DDoS attack is a common two-step DDoS attack in which the attacker manipulates open DNS servers.

NUMBER OF ATTACKS	927,366
AVAILABLE DEVICES	1,617,024

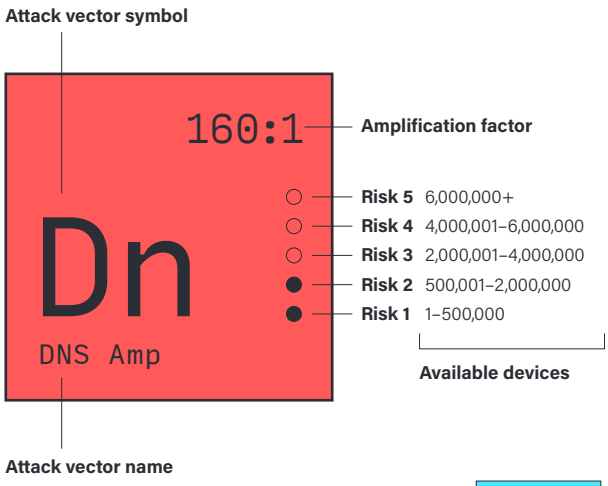

160:1 Dn DNS Amp										Tn TCP NULL
Im ICMP	56.89:1 Cd CLDAP	51,200:1 Mc Memcached							85.9:1 Rd RDP	46.5:1 Tf TFTP
Ta TCP ACK	Ds DNS	25:1 Mq MSSQLRS	3.32:1 St STUN	1,000:1 Ch Chargen	24:1 Di DHCPDiscover	Iv IPv4 Protocol 0	33.9:1 Ov OpenVPN	134.24:1 Ri RIPv1	4:1 Ub Ubiquiti	
Tr TCP RST	1:1 Ik ISAKMP/IKE	556.9:1 Np NTP	35.5:1 Ar ARMS	5.7:1 Ci Citrix-ICA	Ht HTML5	5.6:1 Jk Jenkins	4.68:1 Pm PMSSDP	29:1 Rc rpcbind/portmap	2,464:1 Un Unreal-Tournament	
Ts TCP SYN	13.5:1 Lt L2TP	880:1 Sn SNMP	120:1 Bc BACnet	34:1 Cp COAP	In IP NULL	★ Mh MBHTTP	140.3:1 Qd QOTD	30.7:1 Se Sentinel	14:1 Ve VSE	
Tk TCP SYN/ACK	4.35:1 Md mDNS	30.8:1 Ss SSDP	3.8:1 Bt BitTorrent	37.34:1 Dt D/TLS	1.1:1 Ip IPMI	3:1 Nb NetBIOS	63.9:1 Qk Quake	10:1 Sp SIP	500:1 Wd WS-DD	

500,000+ Attacks | 50,001–500,000 Attacks | 0–50,000 Attacks

HTTP Reflection/Amplification via Abusable Internet Censorship Systems

A largely academic DDoS attack vector thus far, researchers presented a way to amplify a significant amount of attack traffic via abusable internet censorship systems. Until 2021, reflection/amplification attacks were widely believed to be a problem specific to connectionless protocols such as UDP. The USENIX 2021 paper "Weaponizing Middleboxes for TCP Reflected Amplification" proved otherwise. The paper examined a class of middleboxes used by some networks to censor HTTP-based traffic, ultimately showing how middleboxes can reflect and amplify TCP-based application traffic without requiring the sender to first establish a TCP connection. This discovery exposed the susceptibility of these censorship systems to source IP address spoofing attacks, which led to HTTP reflection/amplification attacks.

Vulnerable censorship systems make traffic forwarding or filtering decisions based on the host: header field in an initial client HTTP request. This field typically contains the DNS name the client is attempting to communicate with (e.g., Host: www.netscout.com). If a vulnerable censorship system considers this host name to be prohibited, the request is intercepted, and an HTTP error page is returned. The returned error page is often many times larger than the initial set of address-spoofed packets, and this becomes the amplification component of the attack.

The methods of the attack and the volume of amplification traffic varies. In some cases, practically infinite amplification has been observed due to routing loop configurations of some censorship systems. Vulnerable systems are widely deployed, with tens of millions of IPv4 addresses on the internet exhibiting an application factor of at least two to one. This vulnerability is one of the largest reflection/amplification threats observed to date. Furthermore, the threat is relatively difficult to detect and defend against, because spoofed attack packets can look like ordinary HTTP traffic.

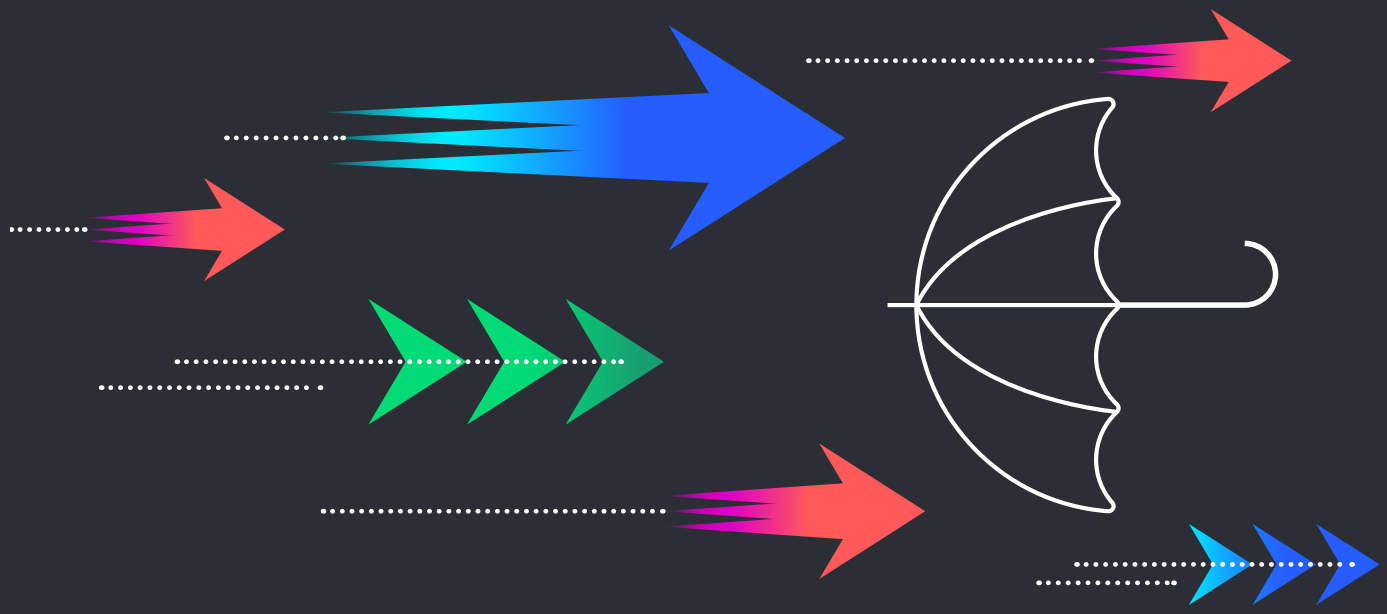

Diving into Direct-Path DDoS Attacks: Fighting Against the Flood

A new era of high-impact DDoS attacks flourished following the introduction of reflection/amplification methodology in 1997. Attackers used to be limited to the bits-per-second (bps) and pps rates directly generated by botnets and customized attack harnesses. Today, however, they punch far above their weight in terms of the amount of amplified attack traffic used against targeted organizations. Worse, easy-to-use DDoS-for-hire services eliminate the technical requirements of launching a massive DDoS attack. Meanwhile, the more mundane direct-path DDoS attacks—such as TCP SYN, ACK, RST, and GRE floods—continue in popularity.

SYN-flood was the most popular DDoS attack vector from 1996 to 2018, when it was overtaken by DNS reflection/amplification. Direct-path DDoS vectors were still employed by attackers either out of habit, randomly, or because of their suitability to task, but reflection/amplification attacks became significantly more prevalent.

But in 2021, reflection/amplification attacks were displaced by direct-path DDoS attacks. This change in trajectory became apparent with the sharp increase in ACK-flood attacks against online credit card processors and other financial services organizations in 1H 2021 and was further supported when SYN floods joined ACK floods as the top two vectors for 2H 2021.

Although there are always myriad factors at work across the DDoS threat landscape, we attribute this increase in direct-path DDoS attacks to the following factors:

OPERATION ANTI-SPOOFING

- Ongoing efforts to implement source-address validation (SAV, commonly referred to as anti-spoofing) by network operators continue to thwart DDoS attackers.

- The ability to spoof source IP addresses is a requirement for launching any type of reflection/amplification DDoS attack. Attack harnesses must be able to forge spoofed attack initiator traffic supposedly sourced from the targeted organization in order to stimulate large, high-impact amplified attack traffic. And although efforts to broadly implement SAV have been ongoing since the early 2000s, it is still not universally deployed—yet.

- As more network operators implement SAV, they deprive attackers of the ability to emit spoofed attack initiator traffic from their networks. This, in turn, limits the breadth of DDoS-for-hire services and bespoke attack infrastructure that can launch reflection/amplification attacks. Although most TCP flooding attacks are spoofed, they are primarily state-exhaustion attacks that are more dependent on packets-per-second throughput rather than bandwidth to negatively impact their targets.

- As the pool of available spoofing-capable bandwidth shrinks, it is often more cost-effective for attackers to launch larger numbers of smaller-bandwidth attacks—especially because high-bandwidth reflection/amplification attacks often include significant collateral damage, thus attracting the attention of both network operators and law enforcement. That higher degree of scrutiny provides additional motivation for network operators to implement SAV even more broadly, further reducing the increasingly constrained pool of spoofing-capable network capacity available to attackers.

- This isn't meant to imply that direct-path DDoS attacks don't generate considerable negative collateral impact. To the contrary, almost all DDoS attacks are overkill, including direct-path attacks, and can significantly interfere with how unrelated parties conduct online activity. However, due to the high-bandwidth focus of reflection/amplification attacks, their collateral damage footprint tends to be even more wildly disproportionate than most direct-path DDoS attacks.

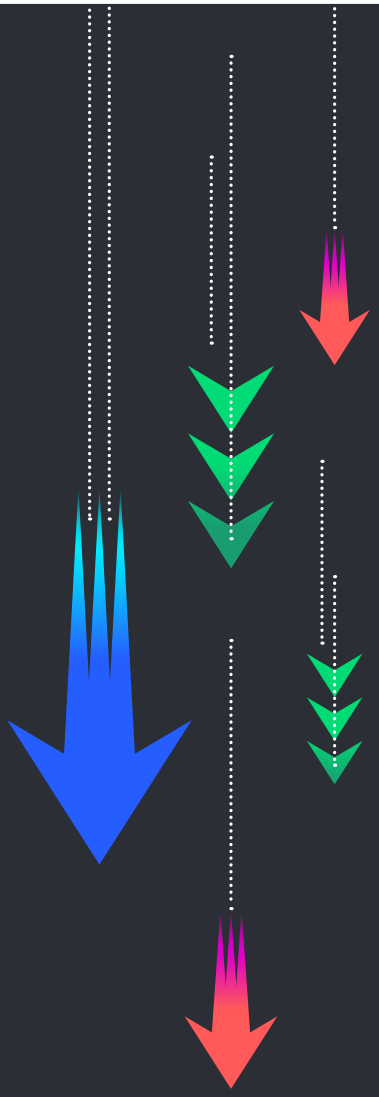

See DDoS-Resistant Architecture for more details.

SERVER-CLASS BOTNET ARMY RECRUITMENT

- The subsumption of server-class nodes into mainstream Mirai botnets means that attackers can launch many simultaneous, moderately scaled direct-path DDoS attacks, while retaining the ability to direct high amounts of attack traffic toward targets on demand. Servers are generally expected to generate significantly more outbound internet traffic than PCs and embedded IoT devices.

- Networks that contain unpatched servers are ripe for takeover and tend to be less closely monitored than networks that are heavily engaged with by the operational security community. As a result, they are more likely to rapidly patch exploitable security vulnerabilities.

- TCP-based direct-path DDoS attacks do not have to be spoofed. When a sufficient number of bots participate in an attack, exhausting state on the attack target can still occur if the defenders are unprepared. Likewise, most application-layer DDoS attacks cannot be spoofed, due to their use of TCP as a transport.

All of these factors drove a marked increase in direct-path DDoS attacks during 2021, and we anticipate that their popularity will continue to grow.

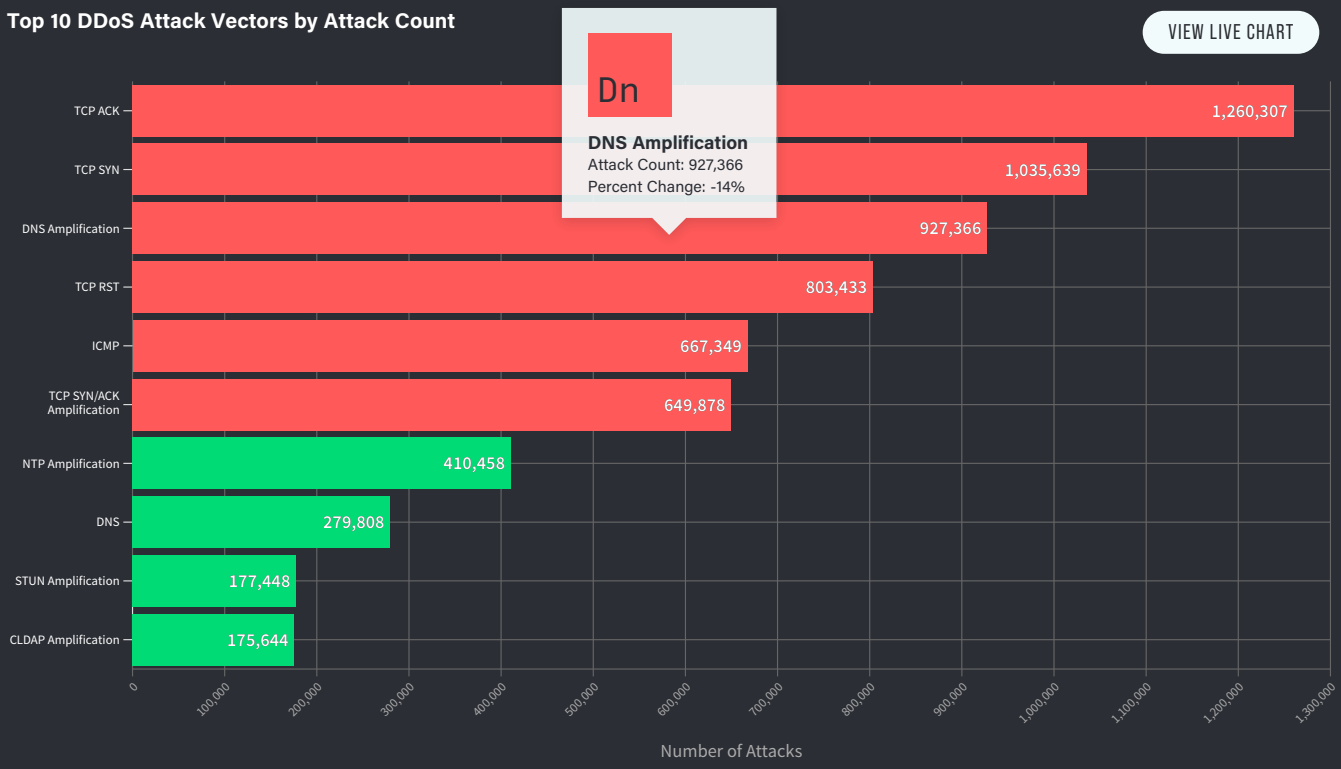

Figure 5: Top 10 DDoS Attack Vectors by Attack Count (Data: Omnis Threat Horizon)

Multivector Attacks and Vector Lifecycles

During the first half of the year, we revealed an omnivector attack in Germany that leveraged 31 different attack vectors, illustrating the upward trend in multivector attacks which we've tracked for more than five years. However, 2H 2021 not only saw a decrease in these attacks for the first time, but that decrease accompanied a significant dip in overall attack numbers and a subsequent decline in some reflection/amplification attacks. This reveals a trend in which adversaries now prefer to use TCP-based floods and botnets for direct-path attacks.

The variability and availability of DDoS attack vectors raises some questions, ultimately spawning an exercise focused on diving into the lifecycle of reflectors/amplifiers over time to reveal patterns of behavior from adversaries that launch such attacks. It's important to note that availability doesn't often equate to DDoS attacks. A good example is the Apple Remote Management (ARM) service. A recent software update from Apple effectively renders this vector moot; however, it doesn't reduce the service's exposure to the internet.

So despite an increasing number of available ARM devices, ARM has seen a significant decrease in usage as an attack vector. From a risk-based approach, vendors and security professionals should seek to both remove from visibility and remediate the exploitable nature of these vectors.

In other cases, however, a decrease in available reflectors/amplifiers has a direct impact on the number, size, and speed of an attack. DNS amplification is one such attack vector that experienced a significant decrease in the number of abusable devices over the last two months of 2021. Incidentally, we observed a 32 percent decrease in DNS amplification attacks. Unfortunately, due to the pervasiveness and constant rotation/addition of new DNS servers, which by their very nature lend themselves to this type of abusability, we anticipate that this trend won't continue. It does, however, serve to illustrate what happens when a significant portion of resources becomes unavailable for adversaries: It accompanies a corresponding decrease in their activity.

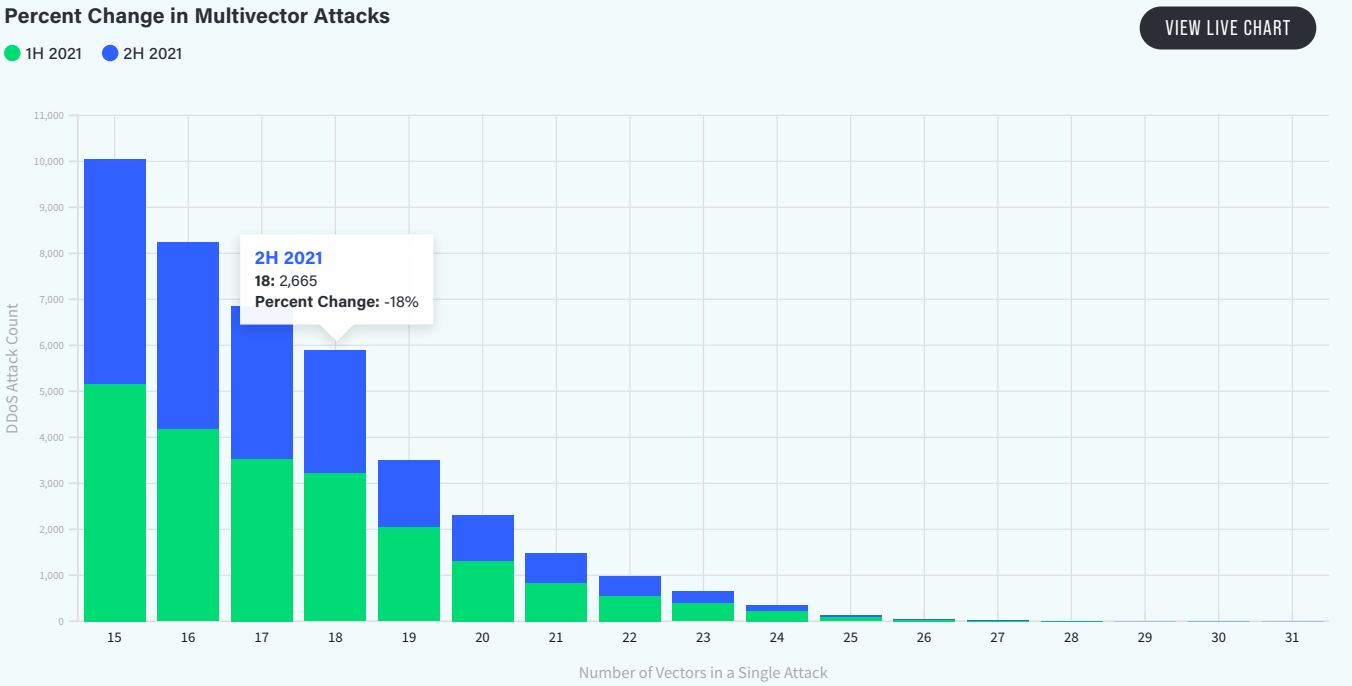

Figure 6: Percent Change in Multivector Attacks (Data: Omnis Threat Horizon)

Timeline of Highly Available Reflectors/Amplifiers by Month (1M–20M)

VIEW LIVE CHART

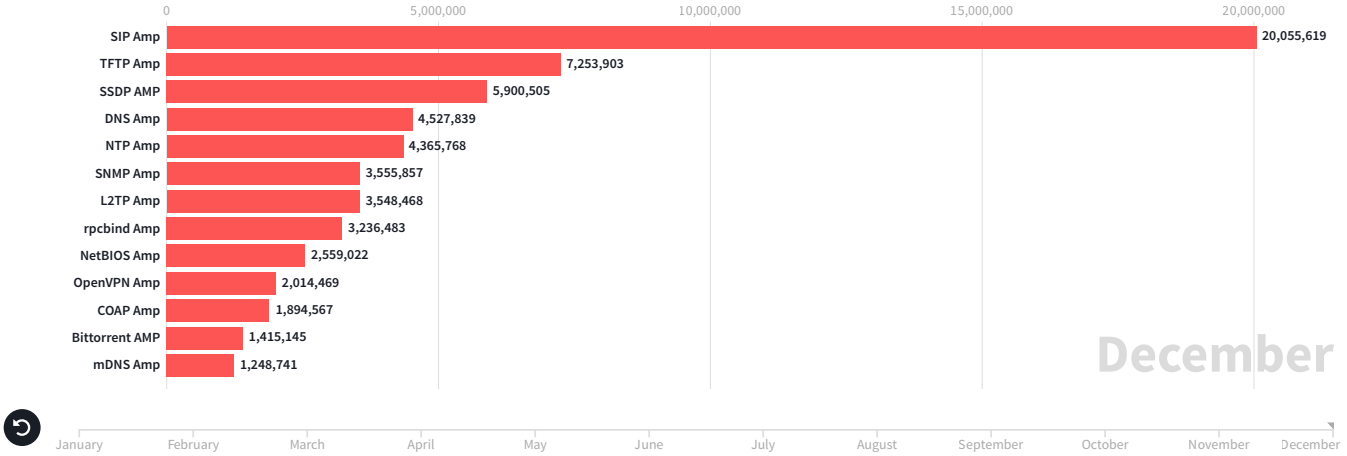

Figure 7: Timeline of Highly Available Reflectors/Amplifiers by Month (1M–20M) (Data: Omnis Threat Horizon)

Timeline of Highly Available Reflectors/Amplifiers by Month (150K–1M)

VIEW LIVE CHART

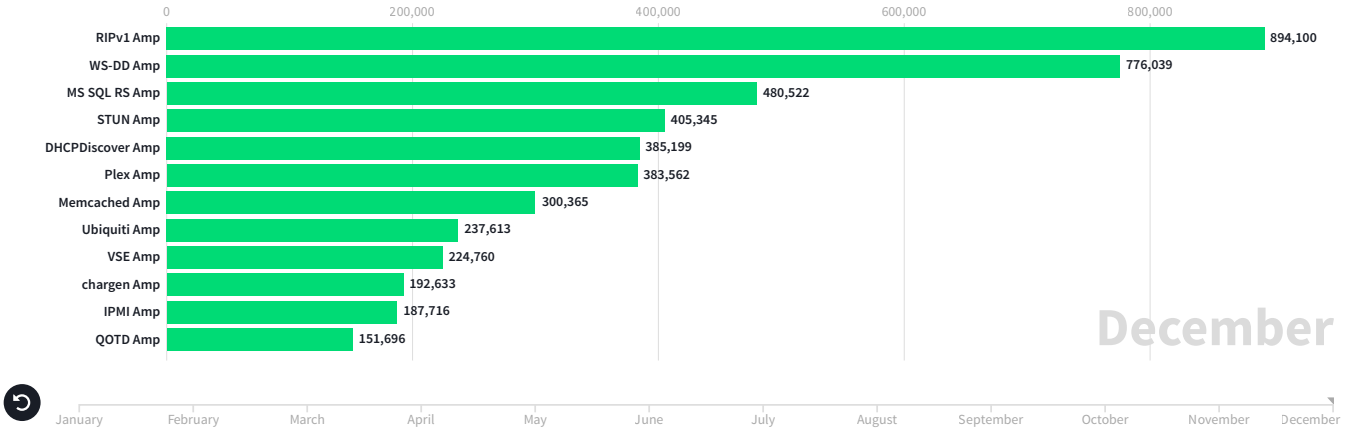

Figure 8: Timeline of Highly Available Reflectors/Amplifiers by Month (150K–1M) (Data: Omnis Threat Horizon)

Timeline of Highly Available Reflectors/Amplifiers by Month (0–150K)

VIEW LIVE CHART

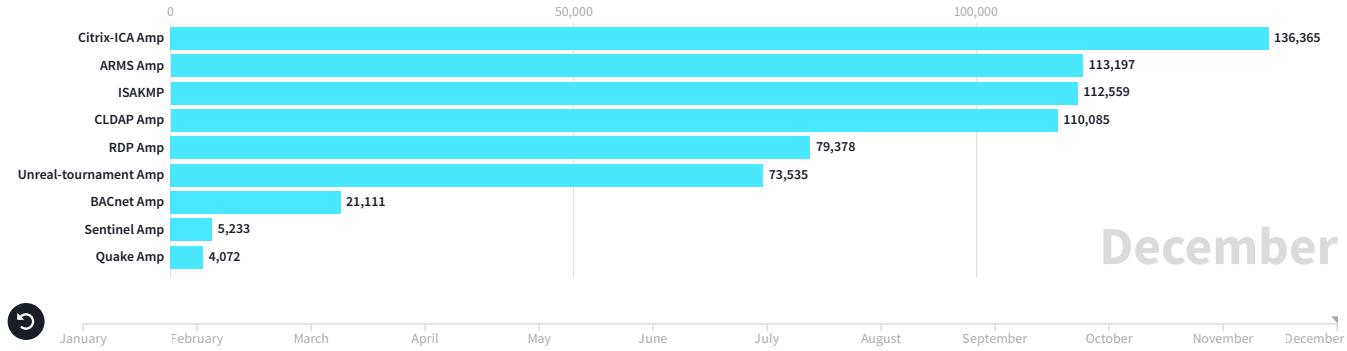

Figure 9: Timeline of Highly Available Reflectors/Amplifiers by Month (0–150K) (Data: Omnis Threat Horizon)

03

Regional DDoS Attack Trends

Across most regions, there was a decrease in attacks, as the global numbers suggest. However, one outlier was an increase of seven percent more attacks in APAC. The past three Threat Intelligence reports chronicle back-to-back declines in attacks for this region. One likely cause of the increase is the geopolitical tensions between China, Hong Kong, and Taiwan—all of which have historically used DDoS as a tool to disrupt activities.

In fact, in June 2019, BBC reported that China was responsible for a powerful DDoS attack that disrupted the Telegram instant messaging platform to curtail communications between protestors in Hong Kong. In December that same year, AT&T security researchers reported that China had resurrected the Great Canon DDoS attack tool to target websites in Hong Kong. Given the propensity for these types of attacks between APAC countries, it's not surprising that cyberattacks have increased with escalated tensions.

Another key geographic point is a four percent increase in attacks against the "all other telecommunications" sector for the EMEA region during a time in which other telecommunications areas (wired and cloud) experienced decreases. VoIP providers, who experienced a significant increase in DDoS extortion attacks, fall into this category.

+7%

Increase in attacks in APAC (despite a decrease in attacks across most regions)

DDoS Attacks by Region

NAMER LATAM EMEA APAC

VIEW LIVE CHART

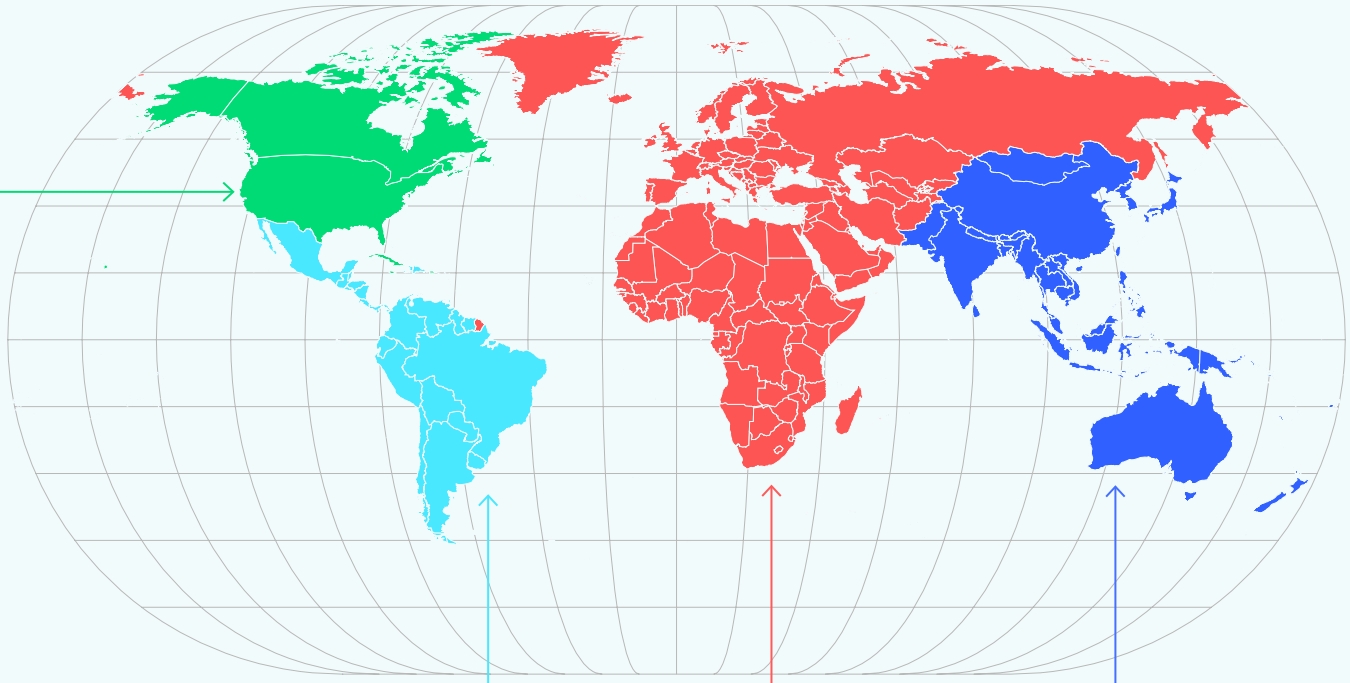

Figure 10: DDoS Attacks by Region (Data: Omnis Threat Horizon)

NAMER

Attack Frequency
962,719 (-24%)

Max Attack Size
554.75 Gbps (-12%)

Max Throughput
269.27 Mpps (+7%)

Average Duration
47 Minutes (+18%)

Top 5 Vectors
- TCP SYN flood
- TCP ACK flood
- DNS amplification
- ICMP flood
- TCP RST flood

LATAM

Attack Frequency
543,534 (-2%)

Max Attack Size
571.06 Gbps (+53%)

Max Throughput
108.58 Mpps (-84%)

Average Duration
70 Minutes (+10%)

Top 5 Vectors
- TCP ACK flood
- DNS amplification
- TCP RST flood
- TCP SYN flood
- ICMP flood

EMEA

Attack Frequency
1,599,665 (-20%)

Max Attack Size
611.58 Gbps (-59%)

Max Throughput
452.52 Mpps (+67%)

Average Duration
47 Minutes (0%)

Top 5 Vectors
- TCP ACK flood
- DNS amplification
- TCP SYN flood
- TCP RST flood
- TCP SYN/ACK amplification

APAC

Attack Frequency
1,267,666 (+7%)

Max Attack Size
522.30 Gbps (-48%)

Max Throughput
222.33 Mpps (-53%)

Average Duration
50 Minutes (-19%)

Top 5 Vectors
- TCP SYN flood
- TCP ACK flood
- TCP RST flood
- DNS amplification
- TCP SYN/ACK amplification

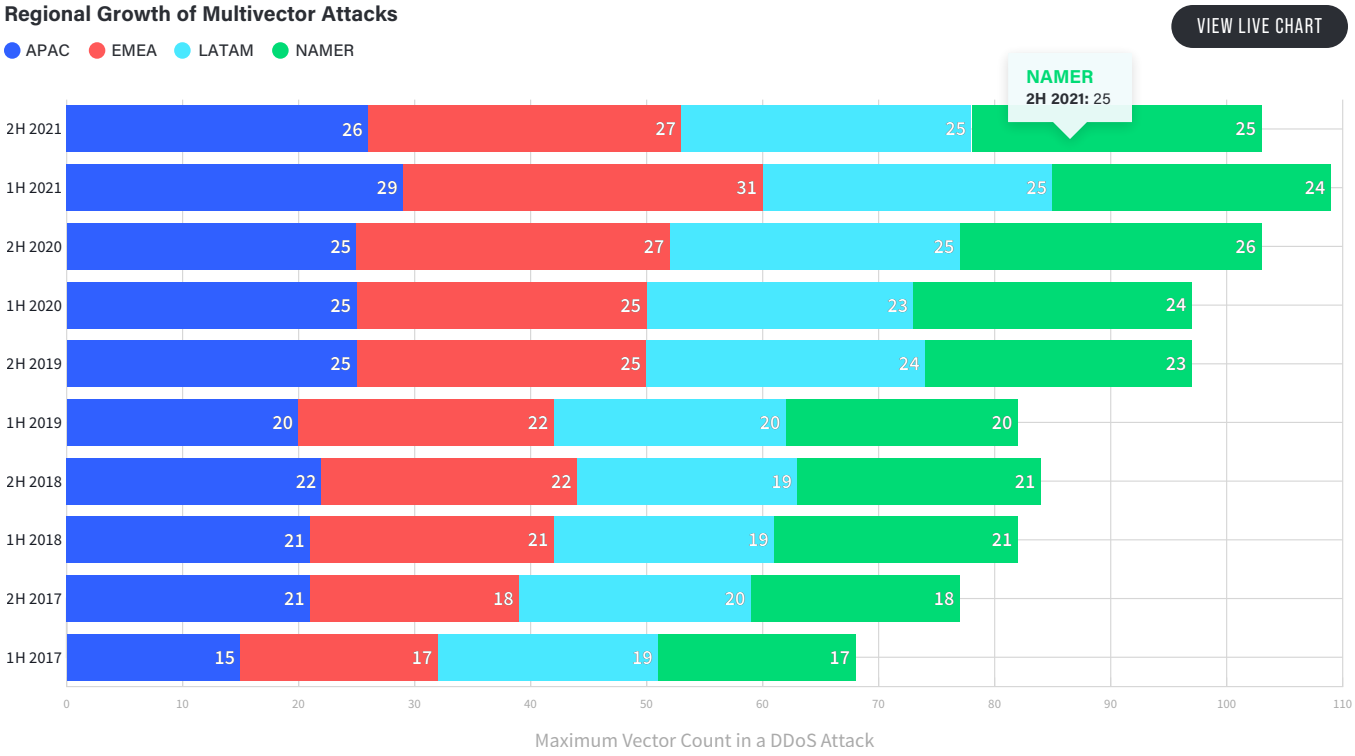

Figure 11: Regional Growth of Multivector Attacks (Data: Omnis Threat Horizon)

Country Snapshots

Read detailed DDoS attack stats across the global threat landscape.

LEARN MORE

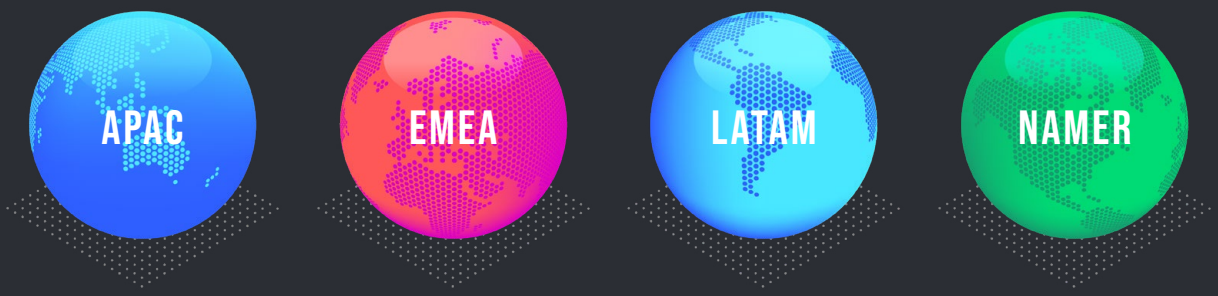

04 Botnet Analsysis

Since 2007, IoT devices have been targeted incessantly by adversaries who try to co-opt them into their botnet armies. Unfortunately, such attacks often are successful because most IoT devices sit behind consumer-grade firewalls—or worse, no firewall at all. In fact, many consumer IoT devices have little to no security, and they're often installed using only default credentials, thereby rolling out a welcome mat for attackers.

Attacks from DDoS botnets on residential networks have limited power due to the fact that most home users lack high-powered bandwidth. The result is that botnet attacks have been carried out via reflection/amplification attacks over direct-path attacks. As is often the case, adversaries are now taking a fresh look at overcoming the limitations of residential devices by using server-class devices to push past the network limitations in home environments. This was first seen with Mēris leveraging HTTP pipelining to launch fast request-per-second (rps) attacks against sites such as Krebs on Security. The success of the attacks quickly led other attackers to piggyback on the vulnerabilities in devices that worked with the Mēris botnet. They did so by leveraging Dvinis to launch more high-powered attacks. Attackers also used Mirai code branches to take advantage of vulnerabilities in GitLab and Confluence servers that essentially recruited them into a server-class botnet army.

ENTERPRISE-LEVEL BANDWIDTH + SERVER HARDWARE = **A new era in high-powered, high-throughput direct-path DDoS attacks**

TOP 3 IoT/LINUX MALWARE FAMILIES

1. **MIRAI**
   Nearly doubled in Q4 with ~203K unique samples (98% increase over 1H 2021)
2. **XOR. DDoS**
   Eclipsed Mirai with ~940K unique samples
3. **GAFGYT**
   ~66K unique samples

Mēris

To perform an after-action review (AAR) of Mēris, it's necessary to look back to when CVE-2018-14847 was first identified in 2018. For three years, this vulnerability enabled adversaries to stealthily compromise MikroTik routers. Those efforts kicked into high gear in June 2021, coinciding with an increase in brute-forcing activity on our honeypot network, as reported in our 1H 2021 Threat Intelligence report. This vulnerability allows attackers to steal unencrypted usernames and passwords of a device after exploitation. As such, system updates failed to mitigate the problem because attackers could still access it via stolen credentials.

The MikroTik platform enabled a retooling of malware, giving attackers access to a much higher level of bandwidth thanks to enterprise deployments of MikroTik devices. It also enabled adversaries to utilize more direct-path DDoS attacks and application-layer attacks.

Mēris Botnet Snapshot

First Seen:

June 2021

Current Active Nodes: ~2,000
Peak Active Nodes: ~4,800
Attacks to Date: ~4,000
Maximum Attack Size: ~337 Gbps
Average Attack Size: ~7 Gbps

Mēris Scanning Details

Mēris nodes continue to bombard our global honeypot with brute-force attempts on RDP, SSH, and Telnet, coinciding with exploitation attempts directly related to the MikroTik router vulnerability.

TCP Port	Count	UDP Port	Count
23	7,234	**123**	486
8291	4,803	**3389**	96
22	3,574	**389**	28
3389	264	**1900**	24
80	113	**3702**	23

Mēris Credential Set

The 1H 2021 Threat Intelligence report highlights a series of MikroTik-specific credential sets that appeared around the time an increase in exploitation of MikroTik routers using CVE-2018-14847 took place. In addition to MikroTik-specific credentials, these username and password combinations were used in an attempt to access our honeypots.

Credential Set	Count
admin:1234	181
root:aquario	160
admin:password	118
admin:123456	100
admin:admin	99

Dvinis

Unlike Mēris, Dvinis-sourced HTTP, and HTTP/S application-layer DDoS attacks don't appear to make use of HTTP pipelining. However, an apparent typo in the attack generators appends an extra "/" character to the end of the Uniform Resource Identifier (URI) targeted in HTTP POST and GET floods. This mistake enables such activity from Dvinis to be tracked.

Additionally, it appears that most of the observed HTTP and HTTP/S DDoS attacks sourced from Dvinis are initiated by an external attack harness and then relayed via the SOCKS4/5 proxy subsystem that's built into compromised MikroTik routers. The HTTP X-Forwarded-For field in captured attack packets includes the source IP addresses of the actual attack infrastructure being used to generate these attacks.

Dvinis Botnet Snapshot

First Seen:

September 2021

Current Active Nodes: ~24,000
Peak Active Nodes: ~24,000
Attacks to Date: ~29,000
Maximum Attack Size: ~463 Gbps
Average Attack Size: ~3 Gbps

Dvinis Scanning Details

As with Mēris, the botted nodes of Dvinis try to propagate across Telnet, SSH, and RDP. Given the massive increase of Dvinis-compromised devices since this activity began, it's clear that spreading attempts have scaled with the increase.

TCP Port	Count	UDP Port	Count
22	18,536	**123**	862
23	14,019	**389**	637
8291	9,963	**69**	457
81	8,526	**137**	434
3389	2,081	**1900**	241

Dvinis Credential Set

Dvinis bots use many of the same top username and password combinations to spread. This is likely due to attempts made to compromise the same kinds of devices by reusing combinations that work. The biggest difference is found in the number of attempts, given that Dvinis has scaled much larger than Mēris.

Credential Set	Count
admin:1234	351
root:aquario	311
root:123456	262
admin:admin	235
admin:12345	232

TRACKING THE DVINIS BOTNET

3,500

Nodes strong in September 2021 when NETSCOUT began tracking

24,000

Nodes strong today (585% increase from September 2021)

GitMirai

Attackers intent on wreaking havoc use GitLab servers to launch terabit-level attacks with incredible throughput. The CVE-2021-22205 vulnerability that was patched in April 2021 allowed botnet commanders to exploit unpatched GitLab servers with a variant of Mirai and the Gitpaste-12 bot, so named by Juniper Networks because it has access to GitLab servers and 12 DDoS attack modules. A report from The Record revealed an attack in the terabit range thus far, and ASERT believes it's only the beginning of bot masters refocusing attention on server-class devices to host their bot code.

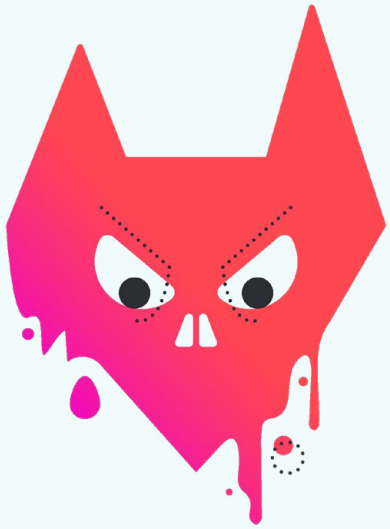

To track the size of this botnet, we examined open ports on GitLab servers (TCP 9418, 80, and 443) and scanned to verify the number of servers. We then correlated the identified servers to our global DDoS attack sensor network to see which had participated in DDoS attacks against our customers in order to ascertain the botnet's level of activity and impact.

GitMirai Botnet Snapshot

First Seen:

November 2021

Current Active Nodes: ~3,800
Peak Active Nodes: ~3,800
Attacks to Date: ~16,000
Maximum Attack Size: ~514 Gbps
Average Attack Size: ~5.4 Gbps

TCP Port	Count	UDP Port	Count
22	58,609	**123**	2,749
2375	1,667	**5060**	1,075
2376	1,441	**389**	659
23	661	**3702**	244
80	434	**3478**	189

GitMirai Credential Set

The following credentials were used most frequently by GitMirai nodes attempting to brute-force our honeypot network.

Credential Set	Count
root:12345678	31
admin:	29
user:user	28
telco:telco	27
default:	26

05

DDoS-Resistant Architecture

DDoS attacks are always present, and adversaries constantly innovate and develop new attack strategies. Nevertheless, it's possible to stop 90 percent of DDoS attacks from being launched with minimal effort by blocking IP address spoofing and controlling inbound traffic.

IP Address Spoofing

When attackers launch reflection/amplification attacks, they often use IP address spoofing, which occurs when a device forges its source address for the purpose of impersonating another device. Doing so forces an unwilling service to send its replies to the victim under attack. There is no practical reason to allow spoofed traffic on the internet; as such, blocking this type of activity has no impact on legitimate traffic.

If IP address spoofing were universally blocked, attackers couldn't launch spoofed DDoS attacks, which would then block all reflection/amplification DDoS attacks. Frustratingly, only 64 percent of autonomous system numbers (ASNs) block IPv4 address spoofing. Likewise, only 78.9 percent of currently announced IPv4 CIDR blocks do so (see the Caida spoofer project). Blocking IP address spoofing is simple to do at the edges of the internet and should be done at the physical edge for each device or at the first routing edge.

It's imperative that corporate networks block IP address spoofing, because attackers look for vulnerable devices inside corporate networks to launch spoofed DDoS attacks. Implementing a access control list (ACL) at the internet-facing edge of the network is a simple process that uses negligible resources, while allowing only legitimate traffic to reach a company network. ISPs should also implement ACLs at the subscriber edges, which allows only inbound traffic originating from subnets allocated to respective customers. This type of control can also be done at the edges between local and regional ISPs, where the regional ISP can control the traffic originating from local ISPs.

ALTHOUGH BLOCKING IP ADDRESS SPOOFING ADDS SOME COMPLEXITY, THE BENEFITS OF DOING SO INCLUDE:

1. Decreasing the frequency and volume of spoofed DDoS attacks
2. Reducing load on ISP infrastructure worldwide
3. Freeing up resources for legitimate internet traffic

Stopping IP Address Spoofing Can Be Done Manually or by Using uRPF

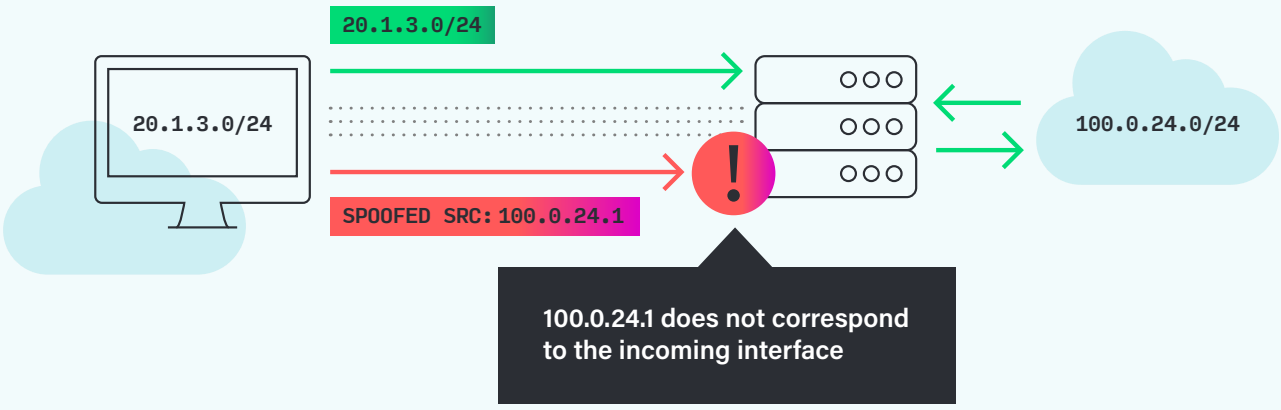

USING ACCESS LISTS TO BLOCK SPOOFING

A simple access control list can be very effective in blocking spoofed packets from internal devices, only allowing legitimate traffic. This can also be automated by using Unicast Reverse Path Forwarding (uRPF).

Controlling Traffic Toward Your Services

Enterprises use the internet for two primary purposes: accessing services/information and providing services/information to others—including things such as web services, SIP services, and DNS services. But no company offers all services to everyone. Because such services are often are specialized, it makes sense to control the kinds of network traffic granted access to them. For example, a web hosting service almost never needs to allow UDP packets toward the service, because all the traffic is inbound on TCP port 80 or 443. Likewise, an authoritative DNS service only needs inbound UPD traffic on port 53 with fallback to TCP port 53 when the topology change Truncate bit is enabled.

By understanding the type of services deployed, it's possible to configure strict access controls, thereby effectively blocking the majority of DDoS attacks with minimal effort. This strategy is especially effective when an attacker launches multivector DDoS attacks, because the majority of attack vectors will be blocked, allowing the security team to focus on attacks that are more serious.

Using Traffic Separation to Defend Against Common DDoS Attack Vectors

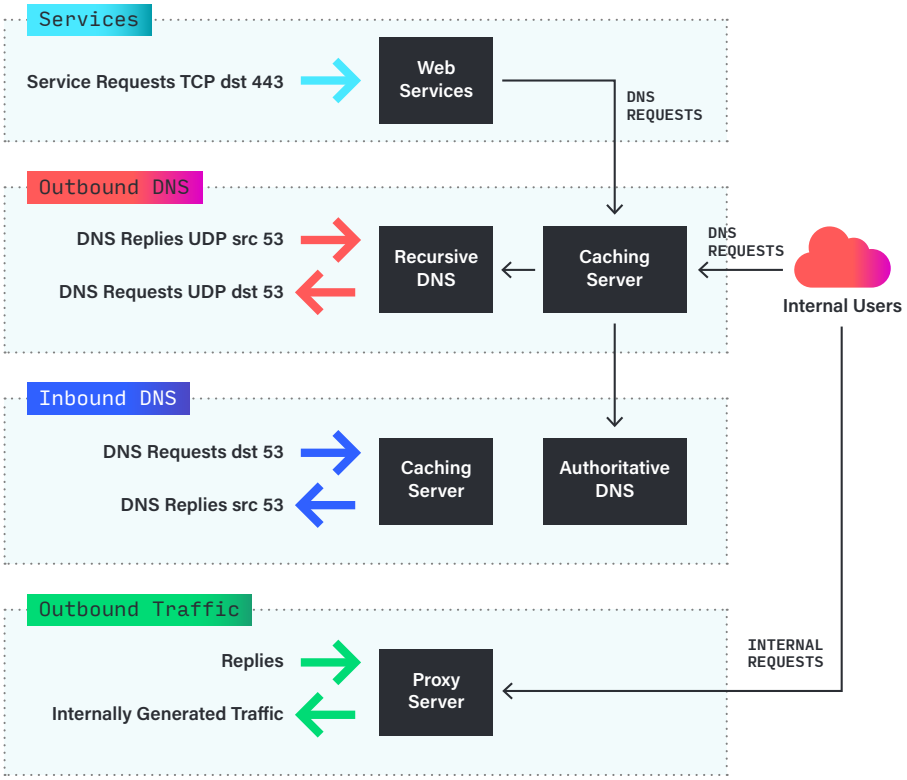

By using traffic separate and strict controls, organizations can block the most common flooding and R/A DDoS type attacks in use today.

THIS ALLOWS ORGANIZATIONS TO:

1. Communicate with external services in a secure way
2. Host services including DNS authoritative servers
3. Simplify the mitigation of TCP SYN type attacks, UDP flooding (including DNS), and DNS R/A attacks

Real-World Examples

1

In 2021, a large service provider that followed these design examples was hit with a massive reflected DDoS attack that attempted to take down its DNS server farm. Without any additional effort, the attack was mitigated by predeployed ACL filters.

2

Likewise, a service provider in Europe that followed these examples faced a similar attack against its authoritative DNS server, and attackers initially were able to disrupt service. It was quickly discovered that a newly deployed edge router was lacking an ACL filter used to block external attacks from source IP addresses. When the ACL was corrected, attack traffic reduced by more than 70 percent, and services were restored.

Summary

By blocking IP address spoofing, implementing best current practices (BCPs), and leveraging intelligent DDoS mitigation solutions (IDMS) such as Arbor Sightline with Sentinel, TMS, AED, and Arbor Cloud, it's possible to fully block or dramatically reduce the impact of DDoS attacks and methodologies like carpet-bombing attacks, TCP-based floods, application-layer attacks, and any other attacks manufactured by adversaries.

06 Conclusion

Despite the drop in overall attack numbers, there's no question that attackers haven't halted their war against corporations, services providers, or connected consumers. In fact, they've become even more entrenched, sharpening their skills with new strategies and mastering techniques to ensure the biggest payday from their extortion efforts. Likewise, attackers continue to add to their tactical playbook, strengthening their botnet armies and running drills using DDoS-for-hire services.

Attackers launched three high-profile DDoS extortion campaigns in 2021—a first-time victory upon which they undoubtedly will continue to build, given that just one of those attacks resulted in at least $9 million in revenue loss. And triple extortion attacks continue to reap massive rewards for attackers, who are constantly innovating and placing new targets in the crosshairs.

In many cases, attackers are targeting organizations and service providers that have been lulled into a false sense of security because they aren't the usual targets. But just because attackers haven't focused as much attention on a particular vertical in the past in no way signals that they won't do so in the future. Indeed, attackers recognize that such companies likely haven't been as stringent in securing networks as they should have been, making them a lucrative target.

So although it's great to see a decrease in attacks to prepandemic days, making security decisions without considering the big picture is a matter of winning the battle but losing the war.

CONTRIBUTORS

Richard Hummel
Roland Dobbins
Chris Conrad
Steinthor Bjarnson
John Kristoff
Denise Culver (Guyer Group)

neustar Security Services

paloalto NETWORKS

UNIT 42 BY PALO ALTO NETWORKS

REVERSINGLABS

ABOUT NETSCOUT

NETSCOUT SYSTEMS, INC. (NASDAQ: NTCT) helps assure digital business services against security, availability, and performance disruptions. Our market and technology leadership stems from combining our patented smart data technology with smart analytics. We provide real-time, pervasive visibility and insights customers need to accelerate and secure their digital transformation. Our Omnis™ cybersecurity advanced threat detection and response platform offers comprehensive network visibility, threat detection, highly contextual investigation, and automated mitigation at the network edge. NETSCOUT nGenius™ service assurance solutions provide real-time, contextual analysis of service, network, and application performance. And Arbor Smart DDoS Protection by NETSCOUT products help protect against attacks that threaten availability and advanced threats that infiltrate networks to steal critical business assets.

To learn more about improving service, network, and application performance in physical or virtual data centers or in the cloud, and how NETSCOUT's security and performance solutions can help you move forward with confidence, visit www.netscout.com or follow @NETSCOUT on Twitter, Facebook, or LinkedIn.

NETSCOUT®

SECR_001_EN-2201 03/2022