



2021–2022

Global Threat Analysis Report

Radware's 2021–2022 threat report reviews the year's most important cybersecurity events and provides detailed insights into the attack activity of 2021. The report leverages intelligence provided by Radware's Threat Intelligence team and network and application attack activity sourced from Radware's cloud and managed services, Radware's Global Deception Network and Radware's Threat Research team.

Contents

| | | | |
|--|-----------|--|-----------|
| Executive Summary | 3 | Ransomware, Now with DDoS Attacks | 37 |
| DDoS Attacks | 3 | Network Scanning and Attack Activity | 38 |
| Geographies and Industries | 5 | Log4Shell | 41 |
| Attack Vectors | 5 | Web Application Attack Activity | 43 |
| Intrusion Attacks | 6 | Security Violations | 44 |
| Web Application Attacks | 7 | Attacking Countries | 46 |
| Unsolicited Network Scanning and Attack Activity | 8 | Attacked Industries | 46 |
| 2021 Retrospective: A Bumpy Ride | 9 | Botnets in Review | 47 |
| Denial-of-Service Attack Activity | 13 | Mozi – The Threat of P2P Botnets | 47 |
| DDoS Attack Trends | 14 | Dark.IoT – Competing for Resources | 47 |
| Attack Sizes | 15 | Mêris – Evolving Tactics | 48 |
| Regions and Industries | 16 | Unsolicited Network Activity | 49 |
| Attack Vectors and Applications | 20 | Most-Scanned and Most-Attacked TCP Ports | 50 |
| Large Attack Vectors | 20 | Most-Scanned and Most-Attacked UDP Ports | 51 |
| Mid-Sized Attack Vectors | 21 | Originating Countries | 52 |
| Micro Floods | 22 | Web Service Attacks | 53 |
| Attack Protocols and Applications | 23 | Top User Agents | 54 |
| Attacks, Attack Vectors and Characterization of Attack Vectors | 26 | Top HTTP Credentials | 56 |
| Attack Vector Characterization | 28 | Top SSH Usernames | 57 |
| Attack Complexity | 34 | References | 58 |
| Record-Breaking DDoS Attacks | 36 | List of Figures and Tables | 63 |
| RDoS and DDoS for Bitcoin on the Rise | 36 | | |

Executive Summary

DDOS ATTACKS

The Radware Cloud DDoS Protection Service mitigated an average of 1,591 attacks per day. The total number of attacks mitigated in 2021 was 580,766. Most distributed denial-of-service (DDoS) activity was concentrated throughout the middle of the year. In the first two weeks of June 2021, the average number of attacks per day was significantly higher and reached a maximum of 9,824 attacks on July 10, 2021. The first half of 2021 had an increasing trend, while the second half had a decreasing trend. The number of attacks mitigated in the first half was almost equal to the number of attacks mitigated in the second half.

The number of blocked malicious events per customer grew 37% from 2020 to 2021. The average attack volume per customer grew 26%. On average, each customer blocked 6.49TB of volume. A DDoS attack in 2021 represented an average volume of 5.69 GB. The largest attack was recorded in Q4 and had a size of 520Gbps.

DDoS Attack Trends in 2021

The number of blocked malicious DDoS events per customer

grew by 37%

Average DDoS attack volumes per customer

increased by 26%

The average volume for large DDoS attacks ranged between

4.6TB and 51.65TB

The average duration of large DDoS attacks ranged between

3.65 hours and 8.72 hours



While less common, several terabit-level attacks were reported in 2021 by large-scale cloud providers. Microsoft Azure reported the largest DDoS attack ever recorded in Q4, with a size of 3.47Tbps. In the same quarter, Microsoft experienced two more attacks above 2.5Tbps.

As businesses migrate critical resources and applications to the public cloud, attackers will have to adapt their tactics and techniques to match the scale of public cloud providers. Enterprises should not immediately be alarmed by these reports of huge attacks. However, they do need to be aware that DDoS attacks are a part of their threat landscape, irrespective of their geography or industry. As such, DDoS mitigation should be part of the protective measures companies implement whenever using or exposing services and applications to the internet.

As bandwidths and resources increase for legitimate businesses, they also increase for threat actors. It is only fair to assume that bad actors can scale as fast and high as their targets. Services hosted in the public cloud will need to consider cloud-scale attacks.

Multiterabit attacks are not necessarily more effective or dangerous than several 100Gbps attacks. In the first few weeks of 2022, during the Twitch Rivals SquidCraft Games event hosted in Andorra, a DDoS attack no larger than 100Gbps interrupted the connectivity of the entire country for hours on end. The attack was performed by an individual or group targeting the event by leveraging a paid subscription to a DDoS-for-hire service.

More concerning is the trend of micro floods and application-level attacks. We noted a slight decline (5%) in the number of large attacks, above 10Gbps, between 2020 and 2021 (see the section “Large Attack Vectors” in “Attack Vectors and Applications”), while attacks smaller than 1Gbps increased by almost 80% (see the section “Micro Floods” in “Attack Vectors and Applications”). Micro floods and slower attacks, such as application-layer attacks, can go undetected and consume resources. Organizations are at risk of having to constantly increase infrastructure resources, such as bandwidth, network and server processing, until the service becomes cost prohibitive. Application-layer attacks typically require more resources to detect them than their network-layer flood counterparts do.

Cloud Adoption

As businesses migrated to public clouds, threat actors adapted their tactics and techniques via “cloud-scale attacks.” Microsoft reported the largest DDoS attack ever recorded, at **3.47Tbps**. It also reported two attacks that were above **2.5Tbps** in Q4 of 2021.



Geographies and Industries

Europe, the Middle East, and Africa (EMEA) and the Americas both blocked 40% of the attack volume in 2021, while the Asia Pacific region blocked 20%. The top attacked industries in 2021 were gaming, retail, government, healthcare, technology and finance. Customers in online commerce and gaming, retail and technology witnessed the largest increase in DoS events and attack volume. Customers in government, healthcare and research and education saw the biggest increase in attack volume. The volume per DoS event for research and education, government and retail saw a severalfold increase between 2020 and 2021. This increase could be indicative of a change in tactics: attacks that were previously random are now being used as part of more targeted and organized campaigns.

Attack Vectors


Radware recorded a slight decline (5%) in the number of attack vectors larger than 10Gbps, but an increase in mid-sized attack vectors of 39% and a steep increase of 79% in the number of micro floods in 2021 compared to 2020.

On average, customers were targeted by 10.8 attack vectors above 1Gbps for every 1,000 attack vectors in Q1 of 2021. This number dropped to 4.93 in Q4 of 2021. Customers found 3.31 attack vectors above 10Gbps per 1,000 attack vectors in Q2 of 2021. Out of every 3,000 attack vectors targeting a customer, fewer than one was above 100Gbps.

In 2021, the most-often-leveraged amplification protocols were NTP, DNS and SSDP. NTP is also the second top-scanned UDP port in Radware's Global Deception Network. Memcached, LDAP, SSDP, SNMP and mDNS, all popular DDoS reflection and amplification protocols, are in the top 10 most-scanned UDP ports recorded by the deception network.

Micro Floods and Application-Layer DDoS Attacks

Micro floods and slower attacks, such as application-layer attacks, can go undetected and consume resources. The number of micro floods **increased by 79%** in 2021 compared to 2020.



The diversity in leveraged attack vectors decreases as the size of the attack vector increases. The average packet size increases with the size of the attack vector. The average attack vector duration also increases as attack vectors become larger and range from a few minutes for micro floods to one hour for attack vectors over 100Gbps. Consequently, the larger attack vectors are also responsible for the largest mitigated volumes in 2021.

Ninety-six percent of the attack vectors recorded in 2021 were smaller than 10Mbps, while the volume generated by those attack vectors represented only 0.3% of the total attack volume in 2021. Sixty percent of the attack volume in 2021 was generated by attack vectors with sizes between 10Gbps and 100Gbps. Attack vectors above 100Mbps represented only 0.8% of all attack vectors recorded in 2021.

TCP attack vectors with throughputs below 10Mbps generated the largest volumes on average and had the longest durations, while UDP attacks were responsible for the highest throughputs and longest durations with attack vectors above 10Mbps. TCP attack vectors were responsible for the highest packet rates and were surpassed in packet rate only by UDP attack vectors for attack vectors larger than 100Gbps.

The average complexity of attacks increased with the size of the attack. The largest number of attack vectors in a single attack was 21 and was an attack between 10Gbps and 100Gbps. Attacks between 10Gbps and 100Gbps had an average duration of 8.72 hours. Attacks below 1Gbps lasted less than an hour, on average.

Intrusion Attacks

Network intrusion attacks consist of easy-to-execute exploits based on known vulnerabilities and range from scanning using open source or commercial tools, information disclosure attempts for reconnaissance, up to path traversal and buffer overflow exploitation attempts that could render a system inoperable or could provide access to sensitive information.

DoS events accounted for one-third of all blocked events in 2021, while intrusions represented two-thirds.

Most of the intrusion activity in 2021 consisted of SIP scanning. The second-most-blocked exploits in 2021 were attempts to exploit a file buffer overflow in Microsoft Internet Explorer through a malformed BMP, a vulnerability that was published in 2004. The third-most-blocked intrusions were Brute Force attempts over SSH.

Log4Shell, arguably the most critical vulnerability of 2021, took the security community by storm in December. Our cloud services detected and blocked more than 800,000 Log4Shell exploits in December and recorded peaks of over 90,000 exploits per day.

1. The OWASP Top 10 is a standard awareness document for developers and web application security. It represents a broad consensus about the most critical security risks to web applications and is published by the OWASP Foundation.

WEB APPLICATION ATTACKS

The number of blocked malicious web application requests grew 88% from 2020 to 2021.

Predictable resource location attacks accounted for almost half of all attacks. In terms of the 2017 OWASP Top 10 application security risks,¹ broken access control and injection attacks represented three-quarters of all attacks recorded in 2021.

Most attacks originated in the United States and Russia, followed by India, the United Kingdom and Germany. The country in which an attack originates typically does not

correspond to the nationality of the threat actor or group. The originating country of the attack will be chosen by the threat actor based on the location of the victim or the country the threat actor wants to see attributed during false flag operations.


The 2021 attack activity was dispersed across an array of industries, with no one vertical standing out. The most attacked industries were banking and finance and SaaS providers, followed by retail and high-tech industries. Manufacturing, government, carrier, transportation, online commerce and gaming, and research and education all had notable levels of activity.



Web Application Attack Trends in 2021


Average blocked malicious web application requests

grew by 88%



Broken access control and injection attacks represented

over 75% web application attacks



UNSOLICITED NETWORK SCANNING AND ATTACK ACTIVITY

The Radware Global Deception Network registered a total of 2.9 billion unsolicited network events and peaked at almost 10 million events in a single day.

A total of 5.7 million unique IPs were recorded in 2021. This represents 0.15% of the available public IPv4 addresses on the internet. The number of unique IPs provides a good measure for the number of malicious hosts and devices involved in scanning and malicious activity on the internet.

SSH was the target of half of all unsolicited TCP activity, followed by IP cams, RDP, VNC and SMB, and only then followed by the most pervasive web application protocols HTTP and HTTPS. Just 2% of the total activity, but still a notable 24 million events, was targeting Redis, an open source, in-memory data structure store used as a database, cache, and message broker for which a remote code execution vulnerability (CVE-2021-32761) was disclosed in July 2021. This allowed an attacker to execute arbitrary code on the target system.

The SIP protocol, used by many VoIP phones and providers, was the most targeted UDP-based service in 2021. VoIP remains critical to organizations to ensure their productivity, and it also made the list of most-targeted services for DDoS attacks in 2021. Vulnerabilities and weak or default passwords in VoIP services allow these to be abused for initial access, spying and moving laterally inside organizations' networks.

NTP, Memcached, LDAP, SSDP/UPnP, SNMP and mDNS were among the most-leveraged protocols for DDoS amplification attacks and comprised over 60% of all unsolicited network activity. These services are continuously scanned and meticulously cataloged by black hat threat actors to abuse for DDoS attacks, and white hat actors assess the risk in the DDoS threat landscape.

The United States was the top attacking country in 2021, generating more than a third of all unsolicited network activity, closely followed by Russia and China, which both were good for about one-fifth of the total activity.

Apache Hadoop YARN was the most eagerly scanned and exploited online service, followed by platforms, routers and Docker APIs powered by Java Enterprise Edition.

Eight out of the top 10 abused credentials leveraged for account takeover (ATO) attempts in online services consisted of the typical weak passwords “admin”, “pass”, “password”, “123456”, “1234”, “1111”, “1234” and empty password, all combined with usernames “admin” or “root”. Almost one-tenth of all the credentials used during online service attacks consisted of “root:icatch99”, a hardcoded credential in digital video recorders (DVRs) from vendor LILIN that was publicly disclosed in March 2020 [\[1\]](#). DVRs are still ubiquitous in the IoT threat landscape, as are the security cameras that feed them.

The credentials “8hYTSUFk:8hYTSUFk” represented 11% of all abused credentials during online service attacks. The exact origins of the credentials are still a bit of a mystery. They were used in an example for passing authentication arguments to a generic web API interaction and exploration module written in Node called Yiff Rewrite [\[2\]](#), an extended wrapper based on the furry API wrapper. The string was also discovered in several malware binaries.

The top usernames leveraged during SSH Brute Force ATO attempts were unsurprisingly “admin”, “user” and “test”. Among the top 10 are also “postgres”, “oracle” and “git”, exposing the most sought for and most likely targeted services for ATO.

2021 Retrospective: A Bumpy Ride

The year started with the aftermath of the supply chain attack on SolarWinds [3], followed closely by a ransomware attack that forced the executives of Colonial Pipeline to shut down their oil distribution [4].

In January, law enforcement and judicial authorities worldwide took part in the arrest of operators behind Emotet, one of the most prolific banking trojans and malware-as-a-service platforms of the past decade [5]. The Netherlands police announced they took control of the Emotet botnet and were able to dismantle its infrastructure and seize data about its customers. In the second half of 2021, however, Emotet reemerged [6] more evasive than before with the help of TrickBot.

In February, cyber specialists of the Security Service of Ukraine took down one of the most active cybercrime groups since the Maze shutdown, Egregor [7]. Egregor earned its reputation after the group successfully breached Barnes & Noble and video game developers Crytek and Ubisoft in October 2020.

“They’re being hacked faster than we can count.”

In March 2021, Microsoft released security updates for Microsoft Exchange Server to patch several vulnerabilities that could be chained together to perform unauthenticated remote command execution on Exchange servers, dubbed ProxyLogon [8]. The Microsoft Threat Intelligence Center (MSTIC) reported discovering active zero-day exploits it attributed with high confidence to HAFNIUM, a China-based threat group [9]. Several multinational corporations disclosed attacks, and ESET reported more than 10 different advanced persistent threat (APT) groups were actively planting web shells in more than 5,000 Exchange servers [10]. At some point, a security consultant at F-Secure Corporation said, “They’re being hacked faster than we can count” [11].

The same month, the Cybersecurity and Infrastructure Security Agency (CISA) released a security advisory to address unauthenticated remote command execution vulnerabilities impacting F5 BIG-IP and BIG-IQ enterprise networking devices [12]. The vulnerability could allow attackers to take full control of a vulnerable system. Several researchers thought it would be good to reverse engineer the F5 Java software patch and post a proof-of-concept exploit, causing a quick uptake in opportunistic mass scanning activity for exposed F5 systems [13].

Also in March, four criminals were arrested in Barcelona for their involvement with FluBot, a mobile banking trojan that infected an estimated 60,000 mobile devices through “smishing” (a form of phishing delivered through SMS). Authorities took down the command and control infrastructure of FluBot, but the malicious campaign was restored within days of the takedown, again with help from TrickBot [14].

In April, Pulse Secure LLC reported [15] a remote code execution vulnerability in its VPN software, with a Common Vulnerability Scoring System (CVSS) score of 10 following a report from Mandiant [16]. The vulnerability allowed threat groups with suspected ties to the Chinese government to bypass authentication and maintain access through web shells that persist across upgrades.

“Open Source Insecurity: Stealthily Introducing Vulnerabilities via Hypocrite Commits”

Also in April, a group of researchers from the University of Minnesota got banned from the Linux codebase, as they were caught submitting a series of malicious code commits that deliberately introduced security vulnerabilities into the official Linux codebase as part of their research activities. Their paper was entitled, “Open Source Insecurity: Stealthily Introducing Vulnerabilities via Hypocrite Commits” [\[17\]](#).

In May, ransom DoS (RDoS) made another entrance with a campaign targeting unprotected assets [\[18\]](#). The attackers chose a new moniker this time: Fancy Lazarus.

“2,000% increase in VPN attacks as organizations embrace a hybrid workplace”

In June, North Korean attackers breached South Korea’s atomic research agency through a VPN exploit [\[19\]](#) and Nuspire released a report [\[20\]](#) outlining an increase of nearly 2,000% in VPN attacks as organizations embrace a hybrid workplace. Also in June, researchers from Agari planted phony passwords on the web and discovered how extremely quick attackers were to test the credentials and observed that most of the accounts were accessed manually and not by bots [\[21\]](#) [\[22\]](#).

Hackers were able to break into EA games through Slack by purchasing a batch of stolen cookies being sold online for US\$10 [\[23\]](#). In the meantime, a remote code execution vulnerability in the Virtual SAN health check plugin of the VMWare vCenter, with a CVSS score of 9.8, was being actively exploited [\[24\]](#). Sonatype caught a new malicious cryptojacking Python package leveraging typosquatting in package names and infiltrating the Python Package Index (PyPI) repository to secretly pull cryptominers on affected systems [\[25\]](#).

In July, Kaseya Limited, an American software company that develops software for managing networks, systems and information technology infrastructure, had its remote monitoring and management software compromised by the ransomware group REvil. On July 13, REvil websites and other infrastructure vanished; and on July 23, Kaseya announced it had received a universal decryptor tool from an unnamed “trusted third party” [\[26\]](#).

The author behind the peer-to-peer Mozi IoT botnet was arrested by law enforcement [\[27\]](#). Cloudflare Inc. reported [\[28\]](#) a 17.2-million-requests-per-second (rps) DDoS attack on a financial industry customer. Researchers at the University of Colorado Boulder and the University of Maryland published an academic research paper that discloses new ways to abuse a flaw in 200 million internet-exposed middleboxes and generate massive DDoS attacks [\[29\]](#).

In August, bot herders behind IPStress published a press release [\[30\]](#) to advertise their capabilities, titled “IPStress offers one of the finest DDoS for hire service.” In the meantime, the botnet Dark.IoT was adding new exploits in record time and was found leveraging a supply chain vulnerability in a Realtek chipset SDK impacting IoT devices from 65 manufacturers within days of its public disclosure [\[31\]](#).

Also in August, UpGuard Inc. discovered 38 million records exposed by misconfigured Microsoft Power Apps [\[32\]](#). In another data leak, a database containing 1.9 million records with names and personal details of individuals on the FBI terrorist watchlist was discovered on a Bahraini server [\[33\]](#).

“DDoS attack cost Bandwidth.com nearly \$12 million”

In September, REvil was back and started attacking new victims and publishing stolen files on its data leak site [34]. Around the same time, a threat actor leveraged REvil’s name and reputation in a ransom letter tied to an RDoS campaign targeting VoIP service providers [35].

Also in September, Yandex N.V. and Qrator Labs reported [36] mitigating a 21.8-million-rps DDoS attack performed by a botnet they dubbed Mēris. The operators behind the LockBit ransomware-as-a-service platform put out a request to hire the operators behind the Mēris botnet.

The Wiz Research Team disclosed [37] a set of vulnerabilities they discovered in the Azure Open Management Infrastructure (OMI) agent, including an unauthenticated remote command execution that allows attackers access as root to Linux virtual instances. In the meantime, Dark.IoT added two new exploits [38], one based on OMIGOD and another based on a supply chain command injection vulnerability impacting IP cameras using firmware by UDP Technology [39]. Dark.IoT also added new defense evasions and increased its payload attack vectors to a total of 13 different DDoS attacks.

On September 16, Matthew Gatrel and Juan Martinez were convicted of federal criminal charges for operating two DDoS-for-hire services: DownThem[.]org and AmpNode[.]com [40].

“The ramifications of this vulnerability are serious and it is a matter of time – likely minutes after the disclosure – before working exploits are publicly available.”

Also in September, VMWare disclosed [41] a new remote code execution vulnerability in vCenter urging its users to patch immediately with the message, [42] “The ramifications of this vulnerability are serious and it is a matter of time – likely minutes after the disclosure – before working exploits are publicly available.”

In October, the RDoS actor claiming to be REvil continued to cause problems for VoIP providers. Later in the year, Bandwidth.com went on record that these DDoS attacks caused a \$700,000 dent in its Q3 revenues and would cost the company close to \$12 million [43] in actual and reputation damages.

The Dutch police sent a final warning to 29 users who paid for illegal DDoS services on the DDoS-as-a-service website MineSearch.rip [44]: “We have registered you in our system and you will now receive a final warning. If similar incidents occur in the future, we will prosecute. In that case, take into account a conviction, criminal record and the loss of your computer and/or laptop.”

The REvil ransomware group shut down its operation for the second time after the group’s new administrator, 0_neday, reported that a third party had compromised its infrastructure [45]. Reuters later reported that the U.S. government was behind the compromise [46].

Also in October, Cisco Talos Intelligence Group reported exploits in the wild leveraging an earlier disclosed vulnerability in Apache HTTP Server [47]. The vulnerability was a path traversal and file disclosure vulnerability that could allow an attacker to map URLs outside of the web server’s document root. The first fix for the vulnerability was insufficient and led to another, new vulnerability that was fixed subsequently.

Two new authentication bypasses in cameras from Dahua Technology were discovered and disclosed [\[48\]](#). Later in October, Best Buy, Home Depot and Lowes dropped Dahua’s Lorex products [\[49\]](#) and took them off the shelves following reports of Dahua products being deemed a threat to U.S. national security by the U.S. Federal Communications Commission [\[50\]](#) and sanctions on Dahua for human rights violations and abuses by the U.S. government [\[51\]](#).

Schreiber Foods in Wisconsin, due to an undisclosed cyber incident, had to close its plant and distribution center for several days in October, leading to a cream cheese shortage in the U.S. [\[52\]](#). Earlier in the year, ransomware attacks on JBS, NEW Cooperative Inc. and Crystal Valley Cooperative [\[53\]](#) demonstrated that the food supply chain is vulnerable, causing the FBI to release a Private Industry Notification [\[54\]](#).

In December, Dark.IoT was found abusing a recently disclosed vulnerability that allows it to hijack TP-Link routers [\[55\]](#). While there are several DDoS botnets actively targeting routers, the Dark.IoT operator must be one of the most active botnet developers of 2021.

“Truly one of the most significant security threats of the past decade”

On December 9, a publicly disclosed Log4j vulnerability took the security community by storm. The vulnerability allowed an unauthenticated attacker to leverage publicly available exploits for remote code execution and was considered the most critical vulnerability of 2021.

Authorities dealt some serious blows to organized crime in 2021, both in the physical and virtual realms. Drug trafficking has seen a record level of arrests in Europe thanks to earlier events where European police hacked encrypted phones used by thousands of criminals. In the virtual world, hacking back and an agreement between the East and the West dealt a serious blow to ransomware. The road got bumpier for ransomware operators and affiliates, and the outlook for Russian crime groups is becoming even darker, now that Russian authorities arrested members of REvil on their own soil.

The year 2021 was the year RDoS confirmed its pervasive presence in the DDoS threat landscape.

“2021 was a bumpy ride” is based on the monthly Radware *Threat Researchers’ Live YouTube streams*.

Denial-of-Service Attack Activity

The number of blocked malicious events per customer, mitigated by Radware's Cloud DDoS Protection Service, grew by 37% in 2021 compared to 2020.

From 2020 to 2021, the average attack volume per customer grew by 26%.

In the first half of 2021, the number of malicious events increased quarter over quarter, a trend that started in the second half of 2020. In the second half of 2021, the number of malicious events reversed the trend and declined each subsequent quarter. For the first time in 2021, the number of malicious events in Q4 were below the number recorded in the same quarter one year earlier.

The volume per customer increased steadily quarter over quarter in 2020. In the beginning of 2021, the volume ramped up significantly. By Q3 of 2021, the volume reached its lowest point recorded in the last two years. The year 2021 ended with volumes per customer comparable to those in Q4 of 2020. While malicious events and volumes mostly increased quarter over quarter in 2020, 2021 was mostly defined by a significant increase in malicious activity in the first half of the year, with declining activity in the second half. The year 2021 ended with similar volumes compared to the end of 2020.

FIGURE 1:
Number of blocked malicious events, normalized per customer

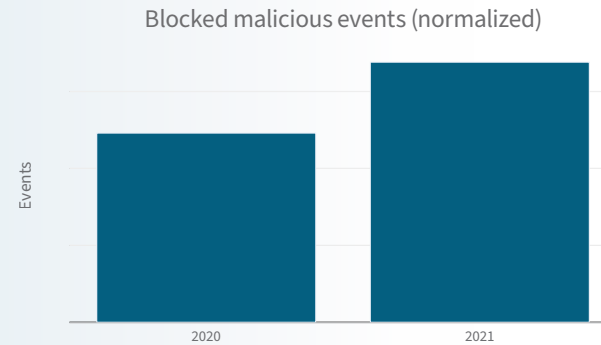


FIGURE 2:
Blocked volume, normalized per customer

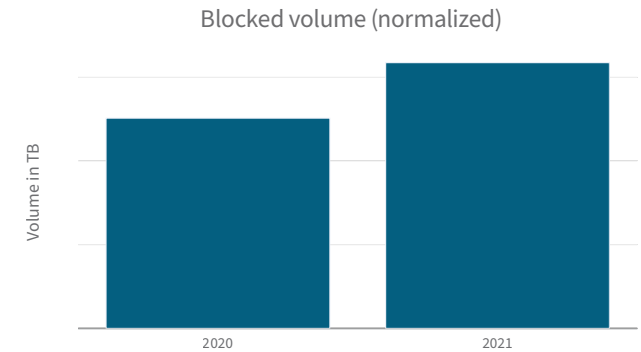


FIGURE 3:
Blocked malicious events, normalized per customer

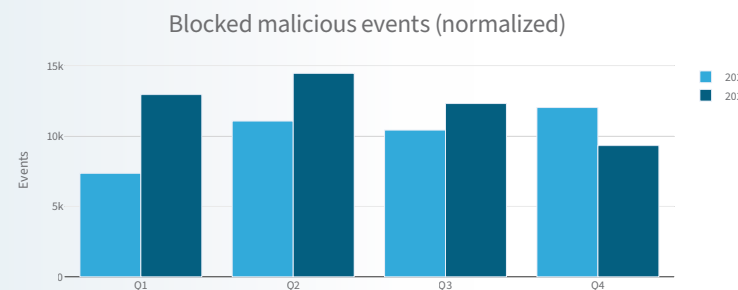
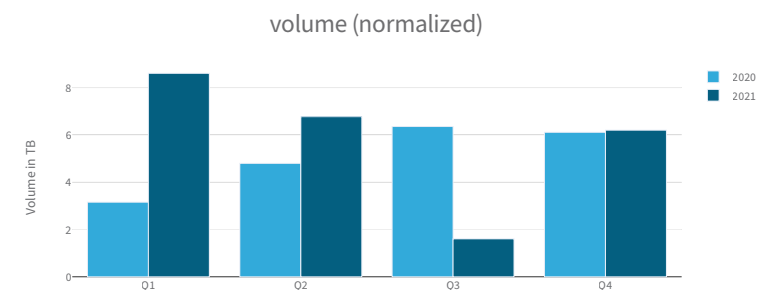


FIGURE 4:
Blocked volume, normalized per customer



DDOS ATTACK TRENDS

The Radware Cloud DDoS Protection Service mitigated an average of 1,591 attacks per day. The total number of attacks mitigated in 2021 amounted to 580,766.

Half of the year’s attacks were mitigated between the start of the year and June 27. During the first two weeks of July, the average number of attacks per day was significantly higher compared to the other days in the year. The number of attacks in one day reached a maximum of 9,824 attacks on July 10.

During the first half of 2021, the number of attacks per attacked customer increased 19% from Q4 of 2020 to Q1 of 2021 and 15% from Q1 to Q2 of 2021. In the second half of 2021, the trend reversed, and there was a decline in the number of attacks per attacked customer, decreasing 11% from Q2 to Q3 of 2021 and 25% from Q3 to Q4 of 2021. By the end of the year, the average number of attacks per attacked customer and the average volume per attacked customer, after reaching record levels of activity, were back to comparable levels to the end of 2020.

FIGURE 5:
Number of DDoS attacks mitigated per quarter




FIGURE 6:
Cumulative sum of DDoS attacks per day throughout 2021

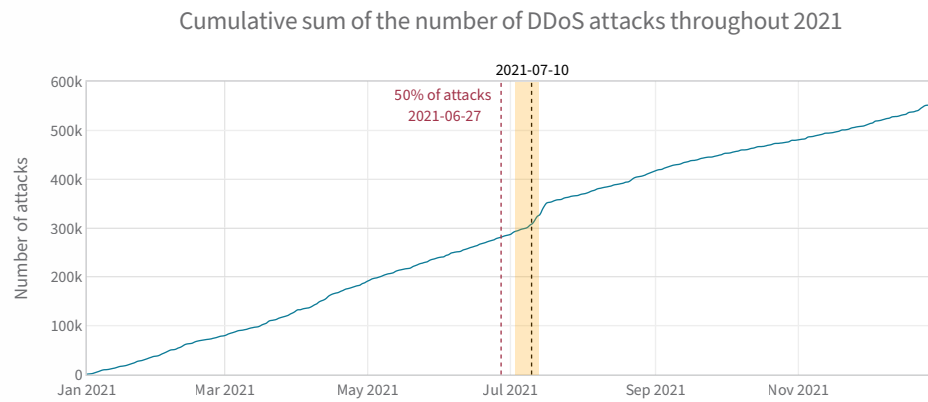
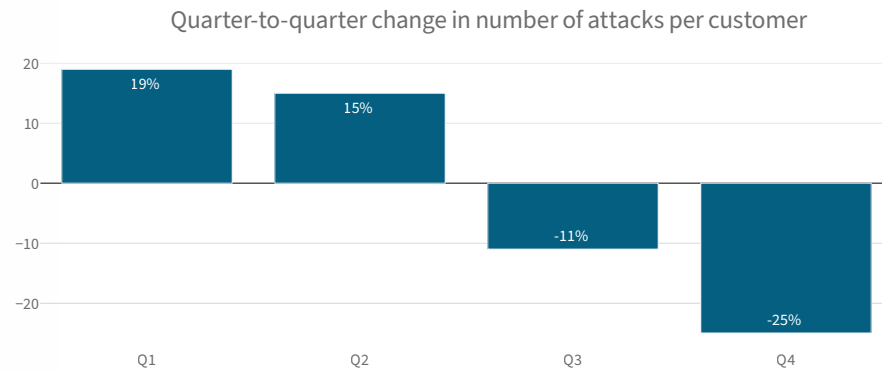


FIGURE 7:
Quarter-to-quarter change in the number of attacks per customer



The average blocked DDoS attack volume per customer was 6.49 TB.

The average volume per DDoS attack was 5.69 GB.

Attack Sizes

The average attack size in 2021, expressed in bits per second (bps), reached its highest average of 162Mbps in Q2. The lowest average attack size was 116Mbps in Q3. The largest attack increased with every quarter, except in Q3. The maximum attack size for 2021 was 520Gbps and was recorded in Q4.

FIGURE 8:
Average DDoS attack volume, normalized per customer

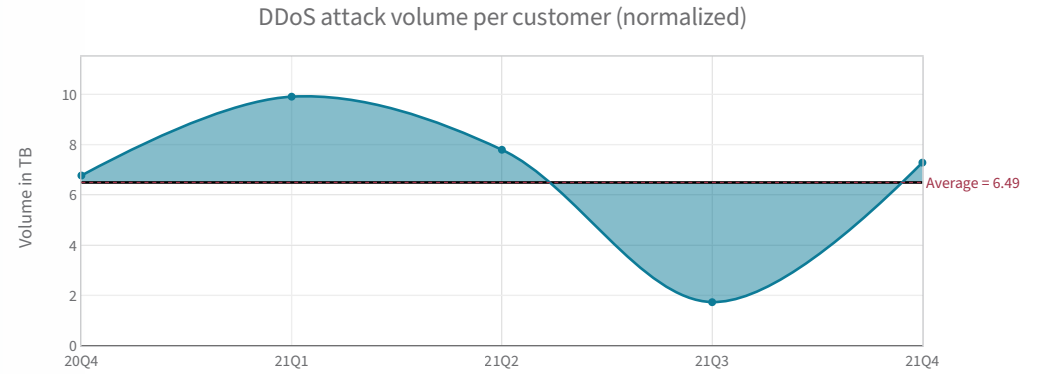


FIGURE 9:
Average volume per DDoS attack, normalized





FIGURE 10:
Average and maximum attack sizes



REGIONS AND INDUSTRIES

In 2020, more than half of the attack volume targeted organizations in EMEA. In 2021, attack volumes were more evenly spread across regions. The Americas and EMEA accounted for 80% of the attack volume, while the Asia-Pacific region blocked 20%.

In terms of normalized attack volume per targeted industry in 2020: gaming, telecom and finance were the most heavily targeted industries. In 2021, there is an almost even spread in normalized attack volume across gaming, retail, government, healthcare, technology, telecom and finance. While mitigated attack volume is just one of several factors to consider when characterizing DDoS attacks, it still provides a good indicator of which industries were most targeted by malicious actors.

FIGURE 11:
Blocked volume per region, normalized, for 2020 and 2021





FIGURE 12:
Blocked volume per industry, normalized, for 2020 and 2021



Comparing the yearly volumes for the top targeted industries provides more insight beyond the relative shifts across industries (see [Figure 12](#)). Gaming, for example, was down from 28.7% to 22.3% of the total attack volume across all industries in 2021 compared to 2020. However, Figure 13 shows that the yearly volume per gaming customer increased by 41% from 2020 to 2021. Increases in mitigated attack volumes were met across most industries, except for telecom and finance.

Table 1 shows the relative change in volume per customer and events per customer by industry, from 2020 to 2021. The highest increases in events per customer were in gaming, technology and retail. Government and research and education had fewer events per customer, but they had a significant increase in mitigated volume per customer, indicating fewer but more-severe attack campaigns.

FIGURE 13:
Yearly volume per industry, normalized

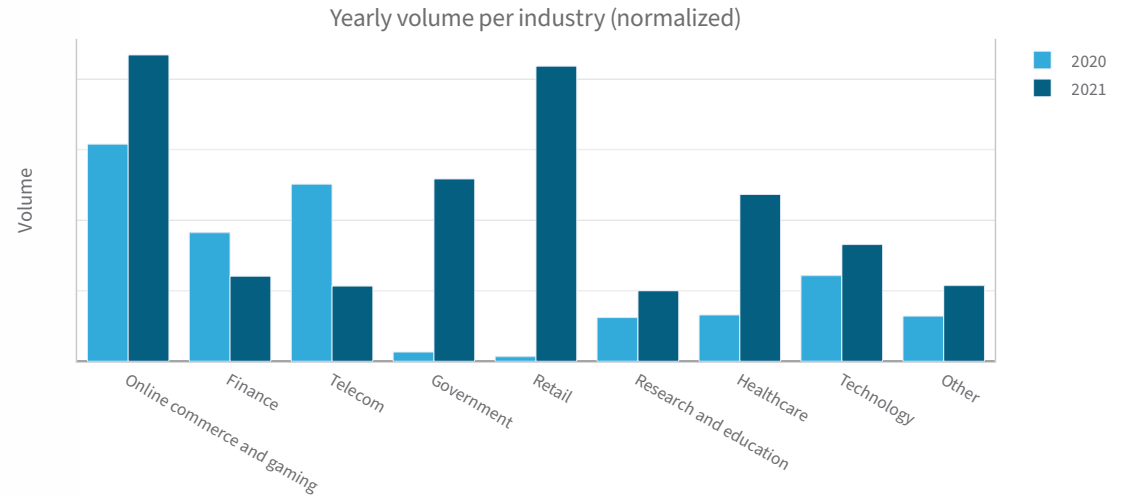


TABLE 1:
Change in volume per customer per industry from 2020 to 2021

| Industry | Change in volume per customer | Change in DoS events per customer |
|----------------------------|-------------------------------|-----------------------------------|
| Retail | +6,288% | +74% |
| Government | +1,881% | -70% |
| Healthcare | +260% | -7% |
| Other | +68% | -47% |
| Research and education | +62% | -96% |
| Online commerce and gaming | +41% | +144% |
| Technology | +36% | +155% |
| Finance | -34% | -28% |
| Telecom | -58% | -53% |

Considering the number of events per customer, gaming and technology were most impacted in 2021; and compared to 2020, they witnessed increases of 144% and 155%, respectively. The number of events in research and education dropped by 96% in 2021, but the mitigated volume per customer increased by 62%.

The volume per event for research and education, government and retail saw a severalfold increase between 2020 and 2021. This increase indicates a potential change in attack patterns from random to more targeted and organized campaigns. Healthcare had a slight increase in volume per event, while technology, telecom and gaming witnessed a slight decrease. Finance and telecom remained mostly unchanged.

FIGURE 14:
Yearly DoS events per industry, normalized

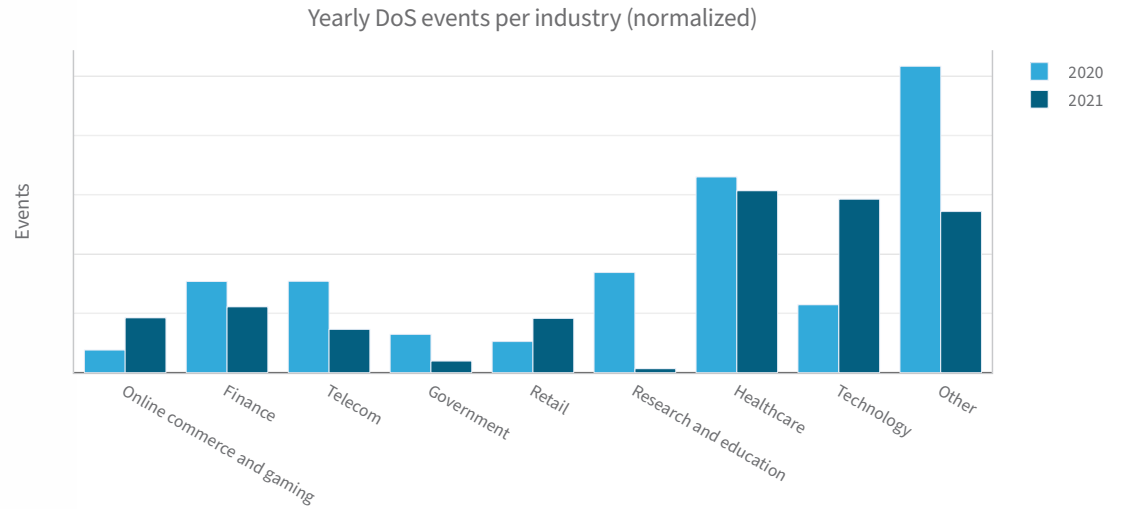


FIGURE 15:
Volume per event by industry




Figure 16 shows the evolution of blocked events and volume per sector over time. Healthcare, for example, had a sharp increase in activity in Q3 of 2020, a trend that continued in Q4 and then slightly declined in Q1 of 2021. The biggest decrease in the number of blocked events for healthcare was in Q2 of 2021, after which the trend slightly decreased in Q3 and Q4. The blocked volume per customer for healthcare, in Figure 17, was severalfold larger in Q1 of 2021 compared to all other quarters.

FIGURE 16:
Quarterly DoS events per industry, normalized





FIGURE 17:
Quarterly blocked volume per industry, normalized



ATTACK VECTORS AND APPLICATIONS

Large Attack Vectors

Radware considers attack vectors above 10Gbps to be large attack vectors. A single large attack vector would be enough to saturate many organizations' headquarters and branches. Not every organization has tens of gigabits-per-second internet links to provide connectivity for on-site employees to cloud-hosted applications or remote access for home workers. Note that this section considers attack vectors. A vector is only one component of an attack. An attack consists of at least one but typically several attack vectors that can be active concurrently or sequentially.

In 2021, the number of attack vectors larger than 10Gbps declined slightly (by 5%). While most of the year, from Q1 until Q3, the number of large attack vectors was higher than the number recorded in the same period in 2020, Q4 of 2020 had a record level of large attack vectors.

FIGURE 18:
Quarterly number of large attacks


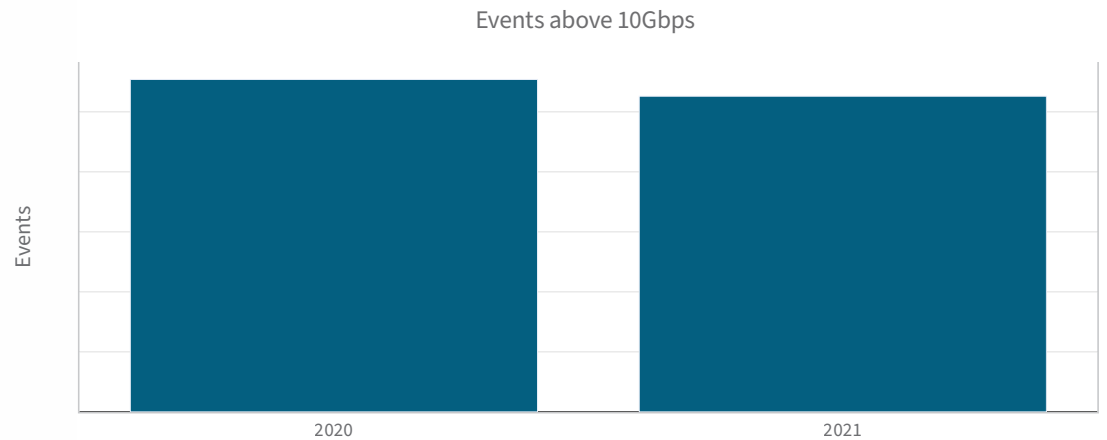


FIGURE 19:
Yearly number of attack vectors larger than 10Gbps



Mid-Sized Attack Vectors

Vectors with throughputs between 1Gbps and 10Gbps are considered mid-sized attack vectors. A single mid-sized attack vector is enough to degrade the quality and experience of internet users and remote workers. Considering that attack traffic comes on top of legitimate traffic, attacks do not always need to reach above the total capacity of the internet connection to degrade the experience of on-premise employees and remote workers.

The number of mid-sized attack vectors in 2021 was consistently higher every quarter compared to 2020. Year over year, the number of mid-sized attack vectors increased 39%.

FIGURE 20:
Quarterly number of mid-sized attacks





FIGURE 21:
Yearly number of mid-sized attacks



Micro Floods

Micro floods, or small attack vectors, are vectors with throughputs below 1Gbps but above 10Mbps to eliminate bias from events that do not qualify as floods. Slower events could be network monitoring probes or discovery scans.

Micro floods do not necessarily impact the user experience. Still, they are enough to become a nuisance when multiple floods are orchestrated concurrently and could force owners to upgrade their internet links or infrastructure to keep a certain level of positive user experience. Micro floods are typically much harder to detect. They are at the bottom of the barrel and cannot be detected using traditional algorithms or techniques that detect larger attack vectors based solely on thresholds.

By combining a large number of micro floods or adding micro floods to a mix of mid-sized and large attack vectors, attackers can increase the complexity of their attack campaigns significantly. Attackers can make mitigation harder by forcing mitigators to constantly adapt their policies.

The number of micro floods increased 79% in 2021 compared to 2020.

FIGURE 22:
Quarterly number of micro flood attacks


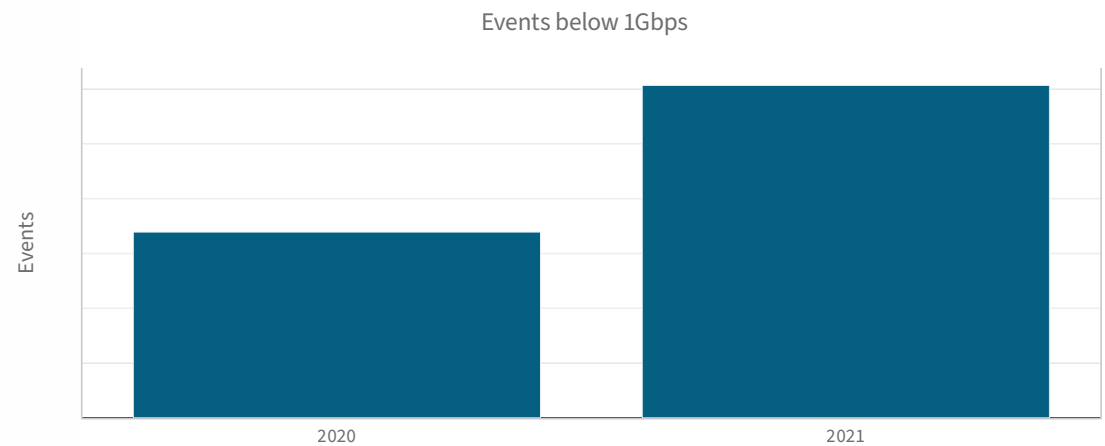


FIGURE 23:
Yearly number of micro flood attacks



On average in Q1, customers protected by Radware were targeted by 10.8 attack vectors of 1Gbps and larger for every 1,000 attack vectors. This number dropped to 4.93 in Q4 of 2021. The probability to be targeted by an attack vector above 10Gbps was between 3.31 and 1.79 per 1,000. The highest probability was in Q2 of 2021. Out of every 3,000 attack vectors targeting a customer, fewer than 1 were above 100Gbps (300 per 1 million).

Attack Protocols and Applications

UDP is by far the most leveraged protocol in DDoS attacks (see Figure 27). Because of its stateless character, UDP allows legitimate services to be abused to send large volumes of unsolicited traffic to victims through reflection and amplification attacks. TCP is harder to trick into sending large volumes because it requires that a connection be established before data is transmitted between client and server. However, TCP still has possibilities for reflection and multiplication of packets. Referring to [29] and [56], TCP can become a very effective amplification vector in terms of packet multiplication and, in some more exceptional cases, a volumetric amplifier.

FIGURE 24:
Number of attacks larger than 1Gbps, normalized per 1,000 attacks

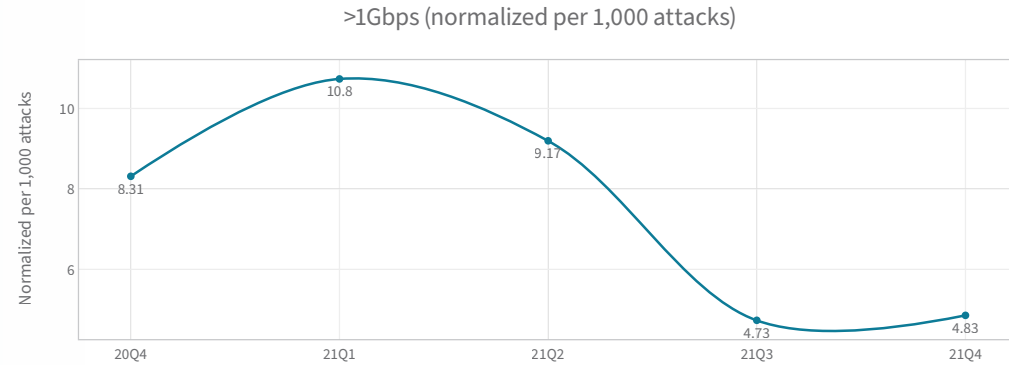


FIGURE 25:
Number of attacks larger than 10Gbps, normalized per 1,000 attacks





FIGURE 26:
Number of attacks larger than 100Gbps, normalized per 1,000 attacks



Attackers primarily favor UDP and leverage many amplification services that are publicly exposed on the internet. If it's UDP and it is exposed to the internet, it can most often be weaponized by DDoS amplification attacks. The motivation to weaponize a specific protocol depends on the amplification factor, or AF (ratio between the size of the request and the reply), and the number of available, exposed services on the internet. A higher AF means a more efficient attack. More exposed services represent a larger total aggregated bandwidth and a higher diversity in source IPs in the attack traffic, making detection (a little) harder.

The main objective behind amplification attacks is saturating a target's connection. Some of the most important and top amplification vectors and their associated maximum amplification factor are listed in Table 2.

FIGURE 27:
Top protocols leveraged by attacks in 2021 (by packets)

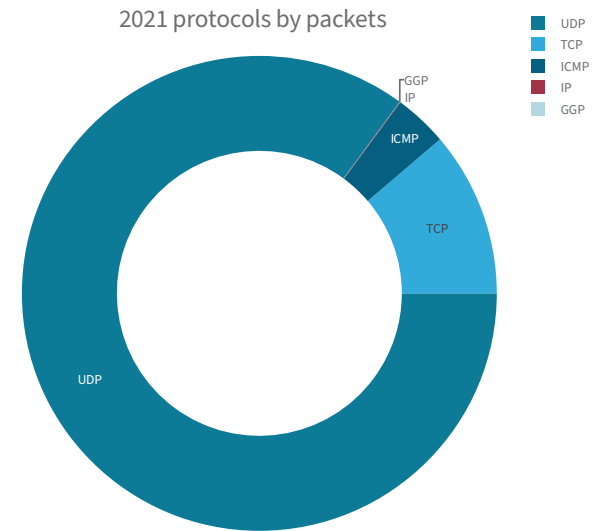


TABLE 2:
DDoS amplification attack vectors

| Amplification Vector | Amplification Factor | Port |
|----------------------|----------------------|--------------------|
| NTP | 500x | UDP/123 |
| DNS | 160x | UDP/53 |
| SSDP | 30x | UDP/1900 |
| Memcached | 50,000x | UDP/11211 |
| Chargen | 1,000x | UDP/19 |
| ARMS | 30x | UDP/3283 |
| CLDAP | 50x | UDP/398 |
| DHCPDISCOVER | 25x | UDP/37810 |
| SNMP | 880x | UDP/161 |
| RDP | 80x | UDP/3389 |
| CoAP | 30x | UDP/5683 |
| mDNS | 5x | UDP/5353 |
| WSD | 500x | UDP/3702, TCP/3702 |
| PMSSDP | 5x | UDP/32410 |

In 2021, the most-often-leveraged protocols were NTP, DNS and SSDP. These protocols represented a significant portion of all amplification attacks in every quarter. Memcached, while present in the first quarter of 2021, is not a consistent attack vector. Memcached services provide amplification ratios that thwart the imagination [57], but they are less likely to be left exposed on the public internet. After their discovery in February 2018, exposed Memcached services have been mitigated, while threat researchers keep scanning the internet for Memcached services that might inadvertently get exposed. An opportunity can arise for threat actors to abuse them, but fortunately, those are typically short lived and rare. Memcached is considered a serious but not persistent threat in the DDoS threat landscape, unlike NTP, DNS, SSDP, Chargen, ARMS and CLDAP.

NTP was the most-used amplification attack vector in 2021, leading every quarter in total dropped packets compared to other amplification vectors (Figure 30) and across the year as the first attack vector behind generic UDP and UDP Frag Floods (Figure 28). NTP is also the second-most-scanned UDP port based on deception network data (Figure 59). Memcached, LDAP, SSDP, SNMP and mDNS are all among the most-scanned UDP ports.

FIGURE 28: Top attack vectors in 2021 (by packets)

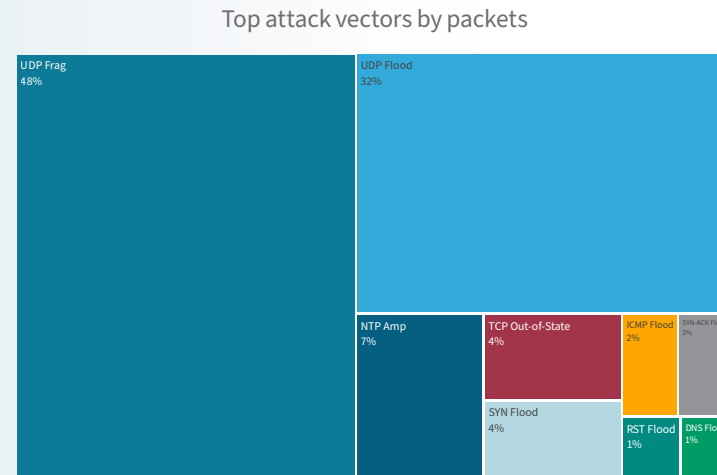


FIGURE 29: Top attacked application protocols in 2021 (by packets)

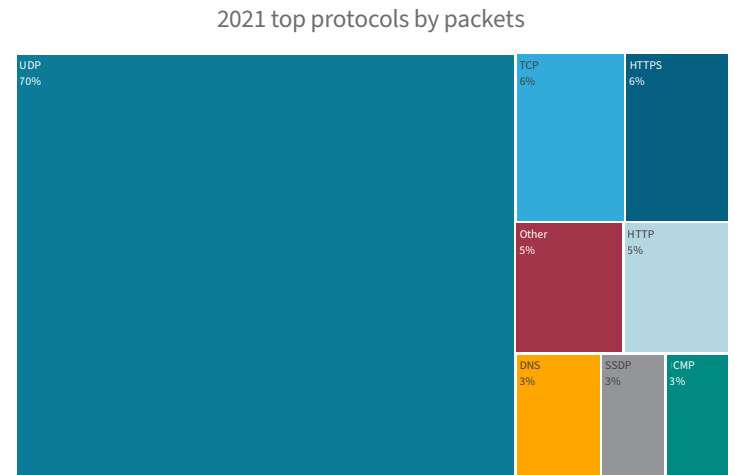



FIGURE 30: Amplification vectors per quarter by volume



ATTACKS, ATTACK VECTORS AND CHARACTERIZATION OF ATTACK VECTORS

The throughput, packet rate and average packet size are characteristic for an attack vector, but they are not independent. There is a linear relationship between throughput in bits per second (bps), the packet rate in packets per second (pps) and the average packet size (\overline{size}) in bytes:

$$bps = (\overline{size} * 8) * pps$$

The relationship between an attack vector’s throughput and rate is mathematically defined by a linear equation, with the average packet size defining the slope and an intercept of zero:

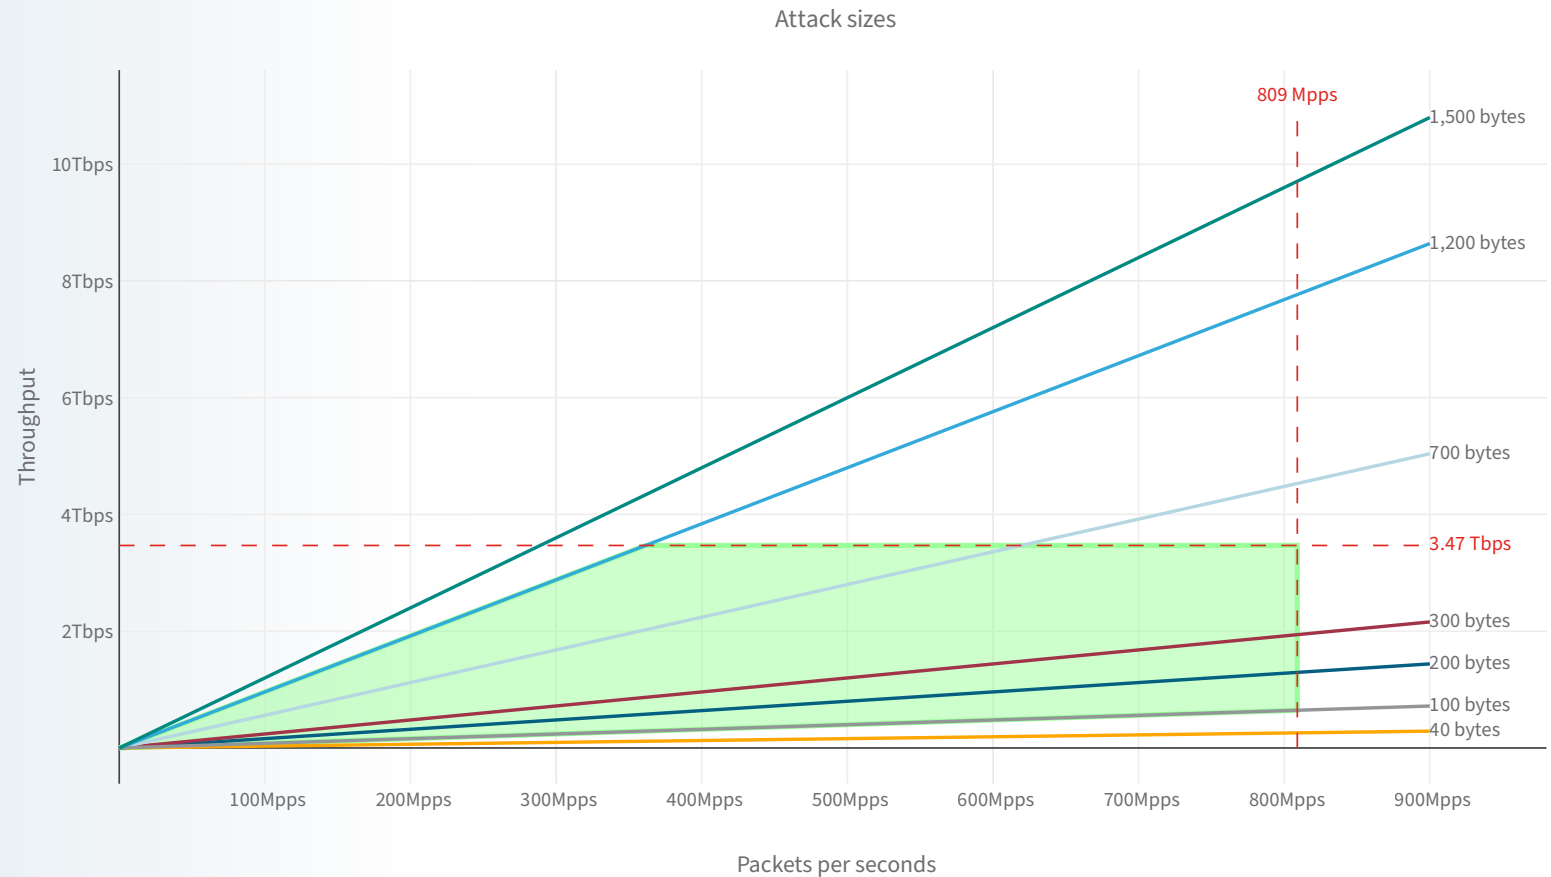
$$y = mx + b$$

where m is the slope and b is the intercept.

The largest DDoS attack throughput ever recorded was 3.47Tbps and was observed in the Azure Cloud during Q4 of 2021 [58]. The highest packet rate ever recorded during a DDoS attack was 809Mpps [59]. Although these numbers were associated with attacks and not attack vectors, we can consider these to be current boundaries for real-world attack vectors.

It is essential to understand the distinction between an attack vector and an attack. A single attack consists of at least one but typically multiple attack vectors. Attack vectors define a certain transfer of packets within a specific

FIGURE 31:
Linear relationship between attack vector throughput and rate in function of packet size



time period. Attack vectors can typically be distinguished by their protocol, source and destination IP addresses and ports. When different random IP addresses or ports are used in an attack, the floods are very similar and can still be considered a single attack vector. A flood with a randomized source IP is called a spoofed flood, while a randomized destination IP within the subnet of the target is a technique referred to as carpet bombing. There is no set or agreed-on definition of what consists of dissimilar attack vectors, and the defining factor for similarity will be imposed by the DDoS detection methods and algorithms.

The most important distinguishing characteristics of an attack vector are:

- **Protocol:** UDP and TCP are most commonly distinguished, as they provide different characteristics that will influence the attack’s impact based on the objective. UDP is typically used in combination with large packet sizes aiming to saturate network connections. These attacks are referred to as volumetric floods. TCP-based attacks typically abuse the statefulness of the protocol and the requirement for devices to keep track of state between individual packets for the duration of the connection. TCP attacks typically leverage smaller packet sizes and try to exhaust resources by hitting systems with high rates of packets, giving them a high administrative workload and a great deal of state to track.
- **Specific services or protocols leveraged for amplification or reflection:** A randomly generated UDP flood is typically distinguished from an amplification attack such as a DNS or NTP amplification.
- **Average packet size:** Except for application-level attacks, packets in a program-generated flood are typically similar in size, independent of whether the flood was generated from multiple devices such as a botnet or a server leveraging amplification and reflection. In the case of amplification and reflection, the total response size of the server might be different between different servers, but the packet size will in all cases approach the maximum allowed packet size if the response is larger than 1,500 bytes on the internet. Note that 1,500 is asymptotic for the average packet size of a flood and represents the maximum transmission unit for Ethernet links.

- **Packets per second and throughput:** The number of packets per second and bits per second are linearly associated through the average packet size. Both measures can be used to characterize an attack. Volumetric attacks are typically expressed in Gbps and Tbps throughput, while high-speed attacks such as SYN floods that target network devices and server resources are mostly characterized by millions of packets per second (Mpps). In this report, the size of the attack refers to the throughput, while the speed of an attack refers to the number of packets per second.
- **Total volume:** This includes the volume generated by the attack during its lifetime.
- **Total packets:** This includes the number of packets generated by the attack during its lifetime.
- **Duration:** This refers to the duration of the attack.

The sophistication of a DDoS attack is partly determined by the number of dissimilar attack vectors leveraged throughout the attack campaign. Depending on the detection, similar attack vectors could be reported individually, but they do not necessarily add to the complexity or sophistication of the attack. For example, a UDP Flood with a random source IP and random source port targeting a specific IP address and specific port is similar to a UDP Flood targeting a single IP address but on another port or randomizing the port or even both the port and target IP addresses within the subnet of the victim’s network range. The latter, which is carpet bombing, is typically leveraged to evade detection systems based on thresholds by spreading the volume as evenly as possible across all IP addresses in the subnet.


In this report, when referring to unique attack vectors, dissimilar vectors are meant. All UDP Floods with different permutations and randomizations of IP addresses and ports are considered a single unique attack vector. Dissimilar vectors would be, for example, UDP Flood, SYN Flood, TCP out of state, DNS amplification and so on.

Attack Vector Characterization

When counting the number of dissimilar attack vectors across all attack vectors recorded in 2021, Figure 32 shows far less diversity in the leveraged attack vectors for larger attack vectors, while smaller attack vectors have a much higher diversity.

Larger attack vectors are typically characterized by a larger average packet size. As noted before, the size (throughput) and speed (packets per second) are linearly associated through the average packet size. To maximize the throughput, the most effective way is to send larger packets. When abusing a service for amplification and reflection, for example, the goal is to trigger a response from the server that is multiple times larger than the request (amplification factor = response size / request size). A larger amplification factor will result in higher efficiency of the attack because fewer requests need to be generated to create an equal-sized volumetric flood. When a response, or any message for that matter, is multiple times the maximum payload of a single packet, the message will be divided across multiple maximum-sized packets and one smaller packet at the end. Ultimately, if the response from an amplification service is larger than a single packet, the average packet size will be higher

FIGURE 32:
Global diversity in attack vectors in function of attack vector size



and defined by the size of the message or the size of the last packet. Asymptotically, amplification attacks will generate an average packet size equal to the maximum packet size of 1,500.

Figure 33 shows the average packet size for different attack vector sizes. The higher the attack vector size, the higher the average packet size. Attack vectors between 10Gbps and 100Gbps had an average packet size of 1,133 bytes, while attack vectors lower than 10Mbps had an average packet size of 103 bytes.

Figure 34 shows the average duration in minutes for different attack vector sizes. The larger attack vectors are associated with longer durations. UDP attack vectors are also considerably longer compared to their TCP counterparts. Attack vectors above 100Gbps average above one hour in duration.

FIGURE 33:
Average packet size
in function of attack
vector size





FIGURE 34:
Average vector
duration per attack
vector size



Note that the overall average duration is just slightly below the UDP duration for most sizes. This is explained by the dominant share of UDP attack vectors, as represented in [Figure 39](#).

Figure 35 and Table 3 show the average volume that attack vectors generate for different attack vector sizes. As noted earlier, the attack bandwidth will depend on the number of packets per second and the average packet size. To generate high bandwidths for large volumetric attacks, a high pps rate and a high average packet size are required. The total volume of an attack vector is determined by the average packet size, the packet rate and the duration of the attack vector. The larger attack vectors are associated with larger average packet sizes ([Figure 33](#)) and longer durations ([Figure 34](#)), so it should not be surprising that they also generated the highest volume per attack vector.

FIGURE 35:
Average volume by attack vector size




TABLE 3:
Average volume per attack vector

| Attack vector size | Average volume per attack vector | Average volume per UDP attack vector | Average volume per TCP attack vector |
|--------------------|----------------------------------|--------------------------------------|--------------------------------------|
| <10Mbps | 4.16MB | 2.76MB | 12.97MB |
| <100Mbps | 658.42MB | 598.91MB | 652.12MB |
| <1Gbps | 11.85GB | 13.56GB | 3.88GB |
| <10Gbps | 201.81GB | 225.98GB | 50.95GB |
| <100Gbps | 3.91TB | 4.13TB | 298.21GB |
| >100Gbps | 25.37TB | 27.67TB | 85.07GB |

Note for attack vector sizes below 100Mbps, the average volume per attack vector is larger for TCP than for UDP. Figure 36 shows this more clearly by charting UDP and TCP average volumes on a logarithmic scale. Attack vectors of similar throughputs and average packet size can vary their volume by changing their packet rates. TCP attack vectors are characterized by smaller throughputs and smaller average packet sizes but higher packet rates (pps). [Figure 34](#) also shows that TCP-based attack vectors are responsible for the longest duration for the smallest vectors. UDP-based attack vectors are responsible for the longest duration for larger attack vectors.

Table 4 shows the average packet rate in pps for different attack vector sizes and by protocol.

FIGURE 36:
Average volume per attack vector, by vector size, log scale




TABLE 4:
Average packet rate per attack vector

| Attack vector size | Average packet rate (pps) | | |
|--------------------|---------------------------|-----------------------|-----------------------|
| | Per attack vector | Per UDP attack vector | Per TCP attack vector |
| <10Mbps | 240 | 132 | 771 |
| <100Mbps | 9,548 | 6,007 | 12,647 |
| <1Gbps | 43,739 | 39,353 | 122,729 |
| <10Gbps | 401,365 | 368,074 | 1,107,449 |
| <100Gbps | 1,066,058 | 1,060,626 | 1,550,099 |
| >100Gbps | 6,145,002 | 6,160,233 | 499,630 |

Note how the average packet rate of TCP attack vectors is higher compared to UDP attack vectors for the same attack vector size. Only for attack vector sizes above 100Gbps does the average UDP packet rate become higher compared to the average TCP packet rate. Figure 37 illustrates this more clearly.

In conclusion:

- UDP attack vectors were leveraged for longer-duration, high-throughput, volumetric attacks and generated large volumes per attack vector.
- TCP attack vectors were leveraged for disrupting network devices and systems by exhausting resources through high packet rates in combination with shorter durations and much smaller volumes per attack vector.

Figure 38 shows the number of attack vectors recorded in 2021 by their attack vector size. Attack vectors smaller than 10Mbps are predominant and represent 96% of the attack vectors recorded in 2021. Attack vectors with a size between 10Mbps and 100Mbps represent 4% of attack vectors, while attack vectors above 100Mbps represent only 0.8%.

FIGURE 37:
Average packet rate per attack vector, by vector size, log scale




FIGURE 38:
Number of attacks per attack vector size




Figure 39 represents the relative shares of UDP and TCP attack vectors for different attack vector sizes. For all sizes, the number of UDP attack vectors represent more than 80%, except for attack vector sizes between 10Mbps and 100Mbps, where the TCP attack vectors are represented by a 43% share. In general, the larger the attack vector size, the higher the share of UDP attack vectors.

UDP attack vectors between 10Gbps and 100Gbps generated 60.3% of the total volume in 2021. Most of the volume is generated by UDP attack vectors, with TCP representing only a negligible part of the total volume. Only 0.3% of the volume in 2021 was generated by attack vectors smaller than 10Mbps, notwithstanding those attack vectors representing 96% of all attack vectors recorded in 2021.

Figure 40 shows how the total volume relates to the attack vector size. For larger attack vectors (higher throughput), the total volume will be larger for an identical duration and proportional to the throughput difference between the attack vectors.

FIGURE 39:
Relative share of UDP and TCP attack vectors





FIGURE 40:
Total attack volume per vector size



ATTACK COMPLEXITY

In this section we are considering attacks, that is, a sequence of events that build up a single attack, consisting of one or more dissimilar attack vectors. An attack is considered more sophisticated or complex when it leverages more dissimilar attack vectors. Attacks that make use of multiple concurrent or attack vectors that change over time attempt to confuse detection and will make mitigation more difficult. Fast shifts and high numbers of concurrent vectors are impossible to mitigate without leveraging automation.

The average complexity of attacks in 2021 increased with the attack size. Since the average number of attack vectors in a single attack can impossibly be smaller than one, smaller attacks exhibit a more isolated character as their average vectors per attack become closer to one. Attacks above 1Gbps average more than two dissimilar attack vectors per attack and double in sophistication for attacks above 10Gbps. Attacks above 100Gbps have, on average, almost seven dissimilar attack vectors.

Attacks between 10Gbps and 100Gbps had the highest average duration, followed by attacks above 100Gbps and those below 10Gbps.

FIGURE 41:
Number of distinct attack vectors per attack in function of attack size


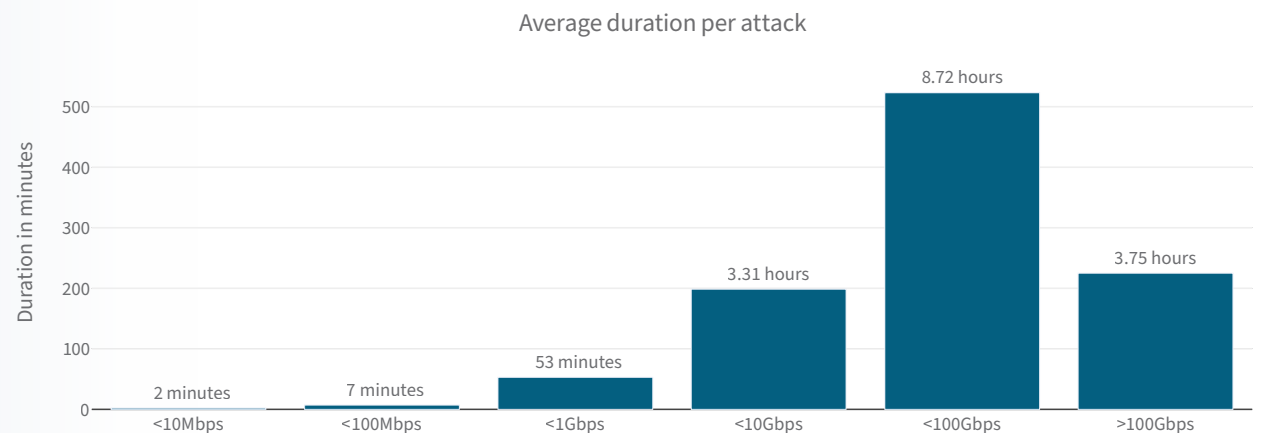


FIGURE 42:
Average attack duration in function of attack size




The average duration for attacks between 10Gbps and 100Gbps was almost nine hours in 2021. The largest attacks were characterized by a shorter duration but still averaged 3.75 hours, close to the average duration of 3.21 hours for attacks between 1Gbps and 10Gbps. Attacks below 1Gbps exhibited a shorter average duration and lasted less than one hour on average.

The volume generated by attacks above 100Gbps averaged at 51.65TB per attack. Attacks between 10Gbps and 100Gbps averaged more than one-tenth lower, at 4.6TB per attack. Attacks smaller than 10Gbps represented on average a volume below 200GB per attack.

The combination of larger throughputs and longer durations contributed to the high averages of the larger attacks.

FIGURE 43:
Average volume per attack in function of attack size



RECORD-BREAKING DDoS ATTACKS

The number of record-breaking DDoS attacks in 2021 was astounding compared to recent years. In total, there were four record-breaking attacks throughout the year. Two of the DDoS attacks were volumetric network-layer attacks, but more impressive, the other two were application-layer DDoS attacks. All these record-breaking attacks occurred within a few months of each other.

In August 2021, Cloudflare reported [\[28\]](#) detecting a world record, a 17.2-million-rps attack originating from 20,000 bots spread across 125 countries. This DDoS attack was an application-layer HTTP attack, a flood of requests designed to consume a significant amount of the server’s resources, therefore causing a DoS condition for legitimate requests. Less than a month later, Qrator Labs reported [\[36\]](#) detecting a similar record-breaking attack in September that generated 21.8 million rps from nearly 56,000 MikroTik devices. Qrator Labs dubbed the attacking entity the Mēris botnet. These massive application-layer DDoS attacks lasted only for roughly 60 seconds, leaving many researchers wondering who and what the objective was behind these attacks.

Following the Mēris attacks, Microsoft reported [\[60\]](#) in October that they detected and mitigated a 2.4Tbps volumetric network-layer DDoS attack targeting an Azure customer in Europe. The attack originated from nearly 70,000 bots in multiple countries in the Asia-Pacific region. As was the case with the other attacks, this one was short lived, with its main burst lasting only 60 seconds.

The fourth and final record-breaking DDoS attack was disclosed by Microsoft in 2022 [\[61\]](#) but occurred in November 2021, a month after Microsoft’s original disclosure of a 2.4Tbps DDoS attack. The attack, targeting an Azure customer in Asia, originated from 10,000 bots located across multiple countries in the Asia-Pacific region and leveraged reflective UDP attack vectors including SSDP, CLDAP, DNS and NTP to achieve a throughput of 3.47Tbps. Like the other attacks, it was short lived and quickly mitigated.

No one has claimed responsibility for these attacks. It wasn’t long ago that hacktivists and DDoS attackers would immediately claim their attacks, or even the attacks of others, via social media. Today, silence seems to be the primary response from threat actors. In some events, ransomware operators such as LockBit have taken to underground forums to ask who launched an attack and whether their botnet or services are available for rent. With this new silent treatment, it has become increasingly more challenging to track criminal activity. But one thing is clear – the threat actors operate increasingly larger DDoS infrastructures.

RDoS AND DDoS FOR BITCOIN ON THE RISE

At the beginning of 2021, Radware published an alert [\[62\]](#) about a ransom denial-of-service (RDoS) group circling back to earlier victims targeted during the summer of 2020. In the new RDoS letters, the group stated that the targeted organizations did not respond to or pay the original ransom demand in 2020 and subsequently would be targeted by a DDoS attack if they were not paid this time. While this event was notable, it was just the beginning of what would become a busy year for RDoS attacks.

In September 2021, another wave of RDoS attacks began targeting VoIP providers worldwide. One of the differing characteristics from this wave of attacks was the group’s name. The threat actors posed as REvil, a notorious ransomware group that had just returned to the threat landscape at the time after completely disappearing following the Kaseya VSA ransomware attack.

The larger RDoS campaigns targeting several VoIP providers such as VoIP.ms, Voipfone, VoIP Unlimited, and Bandwidth.com also sparked concern, as critical infrastructure was heavily impacted by these attacks. This resulted in an industry-wide warning from the Comms Council UK, stating that there was currently a “coordinated extortion-focused international campaign by professional cybercriminals” targeting IP-based

communication services providers. And while generally in the past, RDoS attacks have been considered a low tier threat that's easy to mitigate, one of the victims – Bandwidth.com – expects an impact of \$12 million due to the RDoS attacks.

Another wave of RDoS attacks was observed at the end of 2021, while the VoIP industry was under attack. This campaign targeted multiple email providers such as Runbox, Posteo, and Fastmail. One of the more notable observations from this campaign was that the group calling themselves the “Cursed Patriarch” had more democratic ransom demands of \$4,000. This is similar to the amount requested during RDoS campaigns that initially targeted email providers back in 2015.

RANSOMWARE, NOW WITH DDOS ATTACKS

One of the more problematic developments of 2020 and 2021 was the inclusion of DDoS attacks by ransomware operators. Initially, threat actors exclusively relied on cryptolocking malware to restrict access to user data by encrypting files on systems and devices. Victims were required to pay a ransom in Bitcoin in return for a decryption key. But over time, organizations began training and educating their staff. They refused to pay ransom demands because they had taken precautions and had good backups, forcing threat actors to find new ways to put more pressure on their victims.

In 2019, ransomware groups DoppelPaymer and Maze did just that by doubling down and exfiltrating victim data. If victims decided not to pay the initial ransom because they had backups, they were threatened with the release of sensitive financial, customer or personnel data. Unfortunately, this type of double extortion has become more frequent over the last few years, primarily because threat actors view exfiltration as a backup plan in the event their victims decide not to pay for decryption keys.

Today, there may be close to a dozen or more ransomware groups on the darkweb that leak sensitive files to prove data was compromised. The leak is often amplified when the media picks up on it, and the world soon learns about the latest ransomware victim. In the case of Apple, a journalist wrote an article about what devices were coming out based on leaked content, creating extreme pressure on Apple to protect its intellectual property.

To make matters worse, we now see an added complication to ransomware – a triple extortion threat [\[63\]](#) – exemplified by ransomware groups such as SunCrypt, RagnarLocker, Avaddon, DarkSide, and Yanluowang. Not only does data get encrypted and exfiltrated, but if the victim does not respond to the original threat for payment or the threat of a data leak, attackers may then launch a DDoS attack to bring them back to the negotiation table.

DDoS has traditionally been associated with only one form of extortion: RDoS. Piggybacking DDoS with ransomware is relatively novel. It confirms the increase in sophistication and better organization of cybercriminals. The flourishing underground economy provides threat actors with new and inexpensive ways to rent attack services or keep affiliates on the payroll for additional pressure when required.

So, what can be done to prevent additional layers of pressure from ransomware groups? Not much. Eventually, due to resistance against the threat actors' current tactics, techniques and procedures, along with refusal to pay, ransomware groups will find new ways to pressure their victims into paying. There is too much profit involved for threat actors to walk away. In the beginning, you could survive a ransomware attack with adequate backups. But then, the exfiltration of data made backups alone inadequate. Now, even if you believe you can withstand the exposure of your sensitive data to the public, you must also be able to protect your network against DDoS attacks.


NETWORK SCANNING AND ATTACK ACTIVITY

Not all malicious events targeting internet-exposed assets are DoS attacks. Network-intrusion attacks consist of easy-to-execute exploits based on known vulnerabilities and range from scanning using open source or commercial tools and information disclosure attempts for reconnaissance to path traversal and buffer overflow exploitation attempts that could render a system inoperable or could provide access to sensitive information.

When considering malicious events targeting the same assets and resources, the number of recorded intrusion events is typically larger than the number of DoS attacks. This difference in numbers should, however, not be interpreted as assets having to block more traffic from intrusion than from DoS events. Detection systems are set to report every single intrusion event, amounting to a ratio between packets and events of one to one. DoS events are reported after an attack has ended, and the ratio between packets and events is on average many to one.

DoS events accounted for one-third of all blocked events in 2021. Intrusions represented two-thirds of those events.


FIGURE 44:
*Malicious events by
attack category*



Most intrusion activity in 2021 consisted of SIP² scanning leveraging a tool named SIPVicious. SIPVicious is a set of open-source security tools used to audit SIP-based VoIP systems. It allows discovery of SIP servers, enumeration of SIP extensions, password brute-forcing and scanning for known vulnerabilities. SIP-scanning activity increased considerably in 2021 and moved from second place in 2020 to first place in 2021.

The second-most-blocked exploits in 2021 were attempts to exploit a file buffer overflow in Microsoft Internet Explorer through a malformed BMP image file, a vulnerability that was published in 2004 and tracked as CVE-2004-0566. The exploits moved from eighth place in 2020 to the number-two spot in 2021.

FIGURE 45:
Top network intrusions in 2020 versus 2021



2. SIP, or Session Initiation Protocol, is a protocol that can be used to set up and take down Voice over Internet Protocol (VoIP) calls. It can also be used to send multimedia messages over the internet using PCs and mobile devices.

ZmEu is a vulnerability scanner developed in Romania and was commonly used back in 2012. The scanner searches the internet for phpMyAdmin web services. phpMyAdmin is a portable, web-based, open-source administration tool written in PHP that allows remote administration of MySQL and MariaDB databases. It became one of the most popular MySQL administration tools when MySQL was the preferred back-end database for web applications in the pre–cloud native era. ZmEu is also known for its Brute Force credential-cracking ability through SSH. SSH itself is a top-scanned TCP port for unsolicited network activity, as recorded by Radware’s Global Deception Network (see [Figure 58](#) in the section “Most Scanned and Attacked TCP Ports” in “Unsolicited Network Activity”). ZmEu activity was the number-one intrusion in 2020 and moved to third place in 2021.

| RADWARE ID | CLASSIFICATION | COMMON VULNERABILITIES AND EXPOSURES (CVE) |
|---|-----------------|--|
| SIP-Scanner-SIPVicious | Scanning | – |
| SIPVicious – A SIP information-gathering and scanning tool that detects SIP devices and identifies active extensions on a PBX phone system and the existence of known vulnerabilities | | |
| HTTP-Reply-MS-IE-MalfrmdBMPBO | Buffer Overflow | CVE-2004-0566 |
| Microsoft Internet Explorer Malformed BMP File Buffer Overflow – A vulnerability in the Microsoft Internet Explorer application that could allow a malicious website to execute arbitrary code when a specially crafted BMP file is loaded | | |
| HTTP-MISC-ZMEU-SCANNER | Scanning | – |
| ZmEu – A vulnerability scanner that searches for web servers that are vulnerable to attacks and attempts to guess passwords through Brute Force methods that may lead to DoS | | |

| RADWARE ID | CLASSIFICATION | COMMON VULNERABILITIES AND EXPOSURES (CVE) |
|---|------------------------|--|
| SMTP-MS-Excel-BO | Buffer Overflow | CVE-2007-3890 |
| Microsoft Excel Workspace Index Value Memory Corruption – A Microsoft Excel (2000–2004) buffer overflow attack. Buffer overflow vulnerabilities occur due to programming errors within input validation routines or their absence. Such vulnerabilities can be exploited by diverting the affected application’s path of execution to execute arbitrary code. If exploited successfully, this vulnerability could result in a compromise of the affected system. This buffer overflow can occur by loading a malicious Excel file. In addition, exploitation attempts of a buffer overflow may cause termination of the attacked service, resulting in a potential DoS to the current Excel session. | | |
| DNS-named-version-attempt | Information disclosure | – |
| IQUERY version on named – The BIND named DNS service is vulnerable to an information disclosure attack, allowing an attacker to determine if the server supports IQUERY requests. The information disclosed contains server version information. | | |
| Web-etc/passwd-Dir-Traversal | Information disclosure | CVE-2021-41733 |
| “../etc/passwd” file access with Directory Traversal – Various web servers may be vulnerable to an information disclosure attack that occurs when the web server is misconfigured or contains coding errors that allow access to sensitive files. A recently discovered vulnerability in Apache HTTP Server (CVE-2021-41733) started being actively exploited in the wild in October 2021 [64] . This particular vulnerability was introduced in a recent version of Apache (2.4.49). Users running older versions of Apache are not currently affected. The fix for CVE-2021-41733 in 2.4.50 was found to be insufficient, leading to a second, new vulnerability (CVE-2021-42013) that Apache is now reporting. As a result, version 2.4.51 was released to fully address the issue. | | |

| RADWARE ID | CLASSIFICATION | COMMON VULNERABILITIES AND EXPOSURES (CVE) |
|---|------------------------|--|
| HTTP-Reply-Windows-ANI-LAI-BO | Buffer overflow | CVE-2007-0038 |
| <p>Windows ANI “LoadAnilcon()” – Windows is vulnerable to a buffer overflow attack (MS07-017) that, if exploited successfully, could result in a compromise of the affected system. This buffer overflow occurs due to insufficient checking of “anih” trunks in ANI files in the LoadAnilcon() function. This vulnerability is known to be exploited in the wild by malicious websites. ANI is a graphics file format defined by Microsoft for simple animated icons and cursors on its Windows operating system.</p> | | |
| HTTP-Reply-MS-Excel-IMData-BO | Buffer Overflow | CVE-2007-0027 |
| <p>Microsoft Excel Malformed IMDATA Record Buffer Overflow – Microsoft Excel buffer overflow attack. Exploitation attempts of this vulnerability may potentially result in a DoS to the Excel session. This condition can occur when the crafted Excel media file contains a malformed IMDATA with a zero value as its length. This particular vulnerability can be exploited to terminate the attacked service, resulting in a DoS condition. However, it cannot be used to inject and execute arbitrary code.</p> | | |
| HTTP-Reply-MS-Excel-SubPar-IO | Buffer Overflow | CVE-2011-0097 |
| <p>Microsoft Excel Substream Parsing Integer Overflow – Microsoft Excel is vulnerable to a buffer overflow attack (MS11-021) due to a failure in the code processing 0xA7 and 0x3C-type records in 0x400-type substreams of BIFF files.</p> | | |
| HTTP-APACHE-TOMCAT-GET-OPEN | Information Compromise | CVE-2018-11784 |
| <p>Apache Tomcat HTTP open redirection – A Uniform Resource Identifier (URI) injection vulnerability in Apache Tomcat. The default servlet in Apache Tomcat versions 9.0.0.M1 to 9.0.11, 8.5.0 to 8.5.33 and 7.0.23 to 7.0.90 can be forced to redirect to an arbitrary URI upon presenting a specially crafted URL.</p> | | |

LOG4SHELL

The December 9, 2021, publicly disclosed Log4j vulnerability took the security community by storm. A vulnerability in a pervasively used Java logging library allowing an unauthenticated attacker to leverage publicly available exploits for remote code execution was considered the most critical vulnerability of 2021. Some argued it was the worst vulnerability of the decade.


While Radware assessed [65] the vulnerability to be easy to exploit, we also noted that performing remote code execution was a more involved process and more difficult to achieve while the remote code was executed in the security context of the logging application, which, according to best practices, should run as a limited user. Immediate action was required to close the vulnerability in applications, systems and devices across the globe. The vulnerability could still allow attackers to escalate privileges on compromised systems, move laterally across the network or access back-end databases and information stores accessible by the application.

Scanning and exploit activity was detected and blocked by Radware’s Cloud WAF Service as early as December 9 at 18:00 UTC – only a few hours after disclosure of the vulnerability. By December 10, scanning and exploit activity ran in the range of several thousands of events per day.

As of December 15, a good amount of clear-text activity was blocked by freshly created and deployed Log4Shell signatures in Radware’s network-level DefensePro devices, while exploits leveraging encrypted transport and targeting web applications were detected and blocked by Radware’s AppWall®. AppWall detected Log4Shell exploits on day 1 without requiring specific signatures, because the exploit was possible only by using a Uniform Resource Identifier to a secondary server that was detected as a server-side request forgery violation.

Figure 46 shows the daily number of blocked Log4Shell exploits in December 2021. Peaks of over 90,000 exploits per day were detected. As was the case with other vulnerability scanning activity, a portion of the recorded events and exploits originated from benign actors and organizations performing internet-wide scans to assess risk and inform corporations that may not yet be aware that they are at risk. Bug bounty programs were initiated to motivate vulnerability researchers to discover vulnerable services and organizations. While the numbers are alarming, a portion of the activity can be considered nonmalicious. The size of the nonmalicious portion is unfortunately more difficult to quantify since white-, gray- and black-hat scanners all leveraged similar attack methods. Some of the white-hat scanners were kind enough to identify themselves through web application parameters or user agent strings, but their identifiers were inconsistent at best and did not allow us to make a complete assessment between benign and malicious operations.

FIGURE 46:
Daily blocked Log4Shell activity in Radware's Cloud WAF Service and Cloud DDoS Protection Services



Web Application Attack Activity

The total number of web application transactions blocked by Radware's Cloud WAF Service grew 88% from 2020 to 2021.

During the first three quarters of 2021, the number of blocked transactions steadily increased. In Q4 the number decreased but was still above the quarterly levels recorded in 2020. The activity in every quarter of 2021 was above the activity in the quarters of 2020.

Web application transactions can be blocked by application-specific and custom rules created by a security operations center (SOC). Figure 49 shows the total number of blocked transactions and the share of those transactions that were blocked by signature and behavioral detection modules. Forty-six percent of blocked web transactions were detected and blocked by web application modules based on known malicious behavior and signatures.

To eliminate potential bias introduced by application- and customer-specific security policies, the remainder of this section will consider only attacks detected and blocked based on known malicious behavior, vulnerabilities and exploits.

FIGURE 47:
Yearly blocked web application transactions




FIGURE 48:
Quarterly blocked web application transactions





FIGURE 49:
2021 web application blocked transactions versus attacks blocked by signature



SECURITY VIOLATIONS

The most important security violation – predictable resource location attacks in Figures 50 and 51 – accounted for almost half of all attacks witnessed in 2021. Predictable resource location attacks target hidden content and functionality of web applications. By guessing common names for directories of files, an attack may be able to access resources that were not intended to be exposed. Examples of resources that might be uncovered through Brute Force techniques are old backup and configuration files, web application resources yet to be published and similar. Predictable resource location attack attempts are covered by the 2017 OWASP Top 10 application security risk broken access control, which was ranked fifth in 2017 and moved to first place in the 2021 OWASP Top 10 (see [Figures 52 and 53](#)).

FIGURE 50:
Top security violation types, normalized per customer





FIGURE 51:
Violation types for known web application attacks by quarter



The number-one web application security risk, according to the 2017 OWASP Top 10, is injection attacks, as illustrated by the SQL injection and code injection top violation types in [Figure 51](#). Cross-Site Scripting sits in second place in 2021 and corresponds to the Cross-Site Scripting (A7) OWASP application security risk.

FIGURE 52:
Blocked security violations by 2017 OWASP Top 10 application security risks





FIGURE 53:
Web application attacks per 2017 OWASP Top 10 category



ATTACKING COUNTRIES

Most blocked web security events originated in the United States and Russia. India, the United Kingdom and Germany completed the top five in 2021. It is important to note that the country in which an attack originates does not have to correspond to the nationality of the threat actor or group. Arguably, the country in which the attack originates will most often **not** correspond to the home country of the threat actor. Threat actors leverage anonymizing VPNs, Tor and compromised servers as jump hosts to perform their attacks. The originating country of an attack will be chosen based on the location of the victim or based on the country the threat actor wants to see attributed during false flag operations.

ATTACKED INDUSTRIES

The most attacked industries in 2021 were banking and finance and SaaS providers, which together accounted for over 28% of blocked web application attacks. Retail and high-tech industries were third and fourth, each with almost 12% of blocked web security events, followed by manufacturing (9%), government (6%), carrier (5.9%), transportation (5.2%), online commerce and gaming (4.2%), and research and education (3.6%).

FIGURE 54:
Top attacking countries


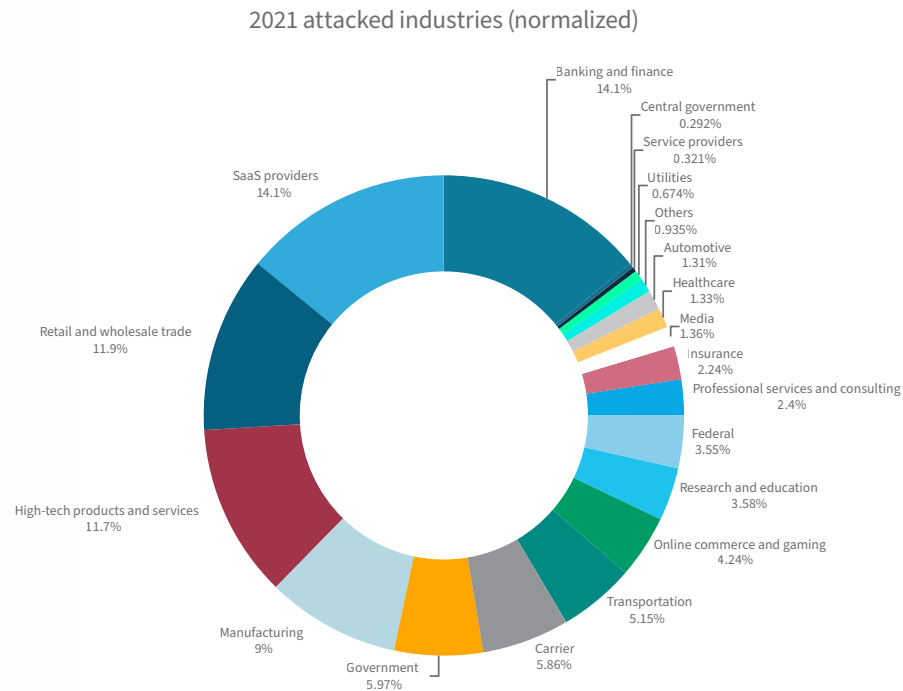


FIGURE 55:
Web application attacks by industry



Botnets in Review

MOZI – THE THREAT OF P2P BOTNETS

Peer-to-Peer (P2P) botnets are nothing new to the threat landscape. In 2019, 360 Netlab discovered [\[66\]](#) a notable P2P botnet dubbed Mozi. This botnet, like Hajime [\[67\]](#), relied on the Distributed Hash Table (DHT) protocol to build and maintain its P2P network. And just like Hajime, it quickly became a noticeable threat triggering alerts in honeypots around the world – at times accounting for over half of all observed malicious IoT traffic.

Mozi spread via telnet Brute Force attacks but was also known to leverage public exploits. The botnet’s primary function was to launch DDoS attacks but could also collect information from compromised devices, download and execute payloads such as mining trojans as well as execute system commands.

In the summer of 2021, 360 Netlab announced that the Mozi botnet was effectively dead due to the arrests of the operators in China. While this was received as great news worldwide, most in the security industry understand how long it takes for botnets like this to fade away. While Mozi will no longer receive updates, it will continue spreading for some time because of its architecture and design. Mozi will eventually disappear only when all targeted network devices are rebooted or updated or the bot is replaced by newer bots.

And just because the operators behind this P2P botnet have been arrested doesn’t mean the security community is clear. In 2020, researchers at Guardicore discovered [\[68\]](#) a new, highly sophisticated, decentralized, and modular botnet called FritzFrog, which shares the same Monero wallet as Mozi. It is expected that this botnet will continue to grow and evolve throughout 2022, setting a new standard for the next generation of botnets.

DARK.IOT – COMPETING FOR RESOURCES

Manga/Dark.IoT [\[31\]](#) [\[38\]](#) was one of the few botnets that gained significant attention throughout 2021. In general, there is nothing special about the malware itself. The botnet is a typical Mirai-based variant that stuck to its primary function – DDoS attacks – and did not diversify its operations to mine crypto or data like Mozi did.

The one thing that did stand out to the security industry was the operator’s ability to quickly evolve and expand their botnet’s capabilities by incorporating recently disclosed exploits into their arsenal. For example, in March 2021, Unit 42 researchers at Palo Alto Networks reported [\[69\]](#) that the operators behind this botnet had leveraged CVE-2021-27561 and CVE-2021-27562 within hours of the vulnerability being disclosed.

The Manga/Dark.IoT campaign in 2021 provided researchers with several opportunities to explore the trials and errors operators face while building and developing a DDoS botnet. One of the most challenging aspects of building a large-scale botnet of any kind is competing with other operators for vulnerable resources. Those that cannot develop or discover exploits on their own are forced to rely on public disclosure. Once a proof of concept is posted, it is a race to be the first operator to leverage the exploit and gather as many vulnerable devices as possible. This process is trial and error, and operators do not always figure out how to properly leverage the vulnerabilities. In contrast, those who do might discover the attempt was not worth their time or effort.

The operators behind the Manga/Dark.IoT botnet attempted to leverage nearly two dozen exploits in 2021. Looking forward to 2022, it is expected that botnet operators will continue to compete for resources by leveraging publicly disclosed vulnerabilities at a high rate.

Defending against these botnet-related attacks will be two pronged. On one side, organizations will need to deploy security solutions designed to detect and mitigate DDoS attacks. On the other side, organizations and the general public will need to manage, update and patch their devices at a much quicker rate moving forward.

MĒRIS – EVOLVING TACTICS

In the third quarter of 2021, Qrator Labs published an article [36] about a recent wave of record-breaking application-layer DDoS attacks. In the report, Qrator Labs attributed these attacks to a new botnet named Mēris (Latvian for “plague”), which is reported to be comprised of more than 250,000 MikroTik devices. The botnet is also said to leverage HTTP pipelining – a process that sends multiple requests over a single connection – to launch short but large-scale DDoS attacks through a network of SOCKS proxies.

To put this in perspective, the largest attack reported from this botnet in 2021 was 21.8 million rps. Not only was this botnet capable of setting records, but it was also used in several RDoS attacks. Because of the botnet’s capabilities and its attention, many operators behind ransomware groups began to seek out those behind Mēris to hire them into their organization.

While questions remain about the botnet, it has been reported that the operators behind Mēris gained unauthenticated, remote access to MikroTik devices via CVE-2018-14847, a vulnerability in the Winbox interface of the MikroTik operating system, RouterOS. Since this vulnerability was patched in 2018, it is believed that the operators are targeting unpatched devices or devices that have been patched but have not updated their default username and password.

The Mēris botnet highlighted the growing and ever-evolving threat landscape around DDoS attacks this year. Going into 2022, we will see the progression and evolution of the threat landscape resulting in more significant DDoS attacks, as threat actors learn to maximize their bots’ resources while staying silent about their work.

Unsolicited Network Activity

Radware’s Global Deception Network consists of a wide range of globally distributed sensors that collect unsolicited traffic and attack attempts. Unsolicited events include DDoS backscatter, spoofed³ and nonspoofed scans and spoofed and nonspoofed attacks.

The difference between deception network events discussed in this section and the web application and DDoS attack events mentioned in previous sections is the unsolicited nature of the events.

Web application and DDoS attack events were collected from services that protect actual services of organizations published and exposed on the internet, backed by real applications and networks. In the latter case, attackers were targeting a particular organization or a known service.

3. IP address spoofing, or IP spoofing, is the crafting of Internet Protocol (IP) packets with a false source IP address for the purpose of impersonating another originating computing system and geolocation. (Source: Wikipedia)

FIGURE 56:
Number of events per month, recorded by Radware’s Global Deception Network

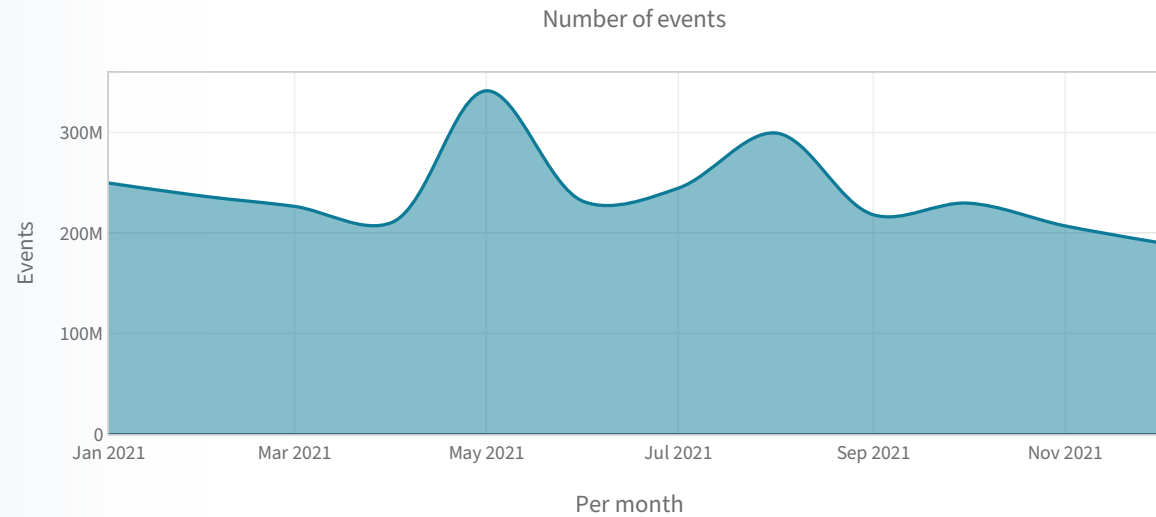



FIGURE 57:
Number of unique IPs per month, registered by Radware’s Global Deception Network



Unsolicited events, as recorded by the deception network, are random acts. The scans or attacks are not targeting known services or a particular organization. The IP addresses of the deception network are not exposed in DNS or used to publish applications or services. No client, agent or device has a legitimate reason to access the sensors in Radware’s Global Deception Network.

The number of events collected by the deception network in 2021 peaked at almost 10 million events in a single day, more than 340 million events per month and a registered a total of 2.9 billion unsolicited events (see [Figure 56](#)).

The number of unique IP addresses provides a measure for the evolution of the number of malicious hosts and devices randomly scanning the internet and exploiting known vulnerabilities. In March 2021, the number of unique IP addresses reached 901,146. A total of 5.7 million unique IPv4 addresses were recorded in 2021, representing 0.15% of the 3.7 billion addresses available for nonreserved use in IPv4 (see [Figure 57](#)).

MOST-SCANNED AND MOST-ATTACKED TCP PORTS

For TCP services, the most scanned and attacked service was SSH on port 22, followed by HTTP on port 8088 (a popular port for IP camera web GUIs), RDP on port 3389, VNC on port 5900, SMB on port 445, and only then came the most pervasive ports 80 (HTTP) and 443 (HTTPS).

Redis (port 6379) is an open-source (Berkeley Software Distribution licensed), in-memory data structure store used as a database, cache, and message broker. In July, a remote code execution vulnerability (CVE-2021-32761) was disclosed due to an integer overflow that affects authenticated client connections on 32-bit versions. A remote attacker can pass specially crafted data to the application, trigger integer overflow and execute arbitrary code on the target system. In April 2020, Trend Micro reported more than 8,000 unsecured Redis instances deployed in public clouds [\[70\]](#).

Telnet on port 23, HTTP on port 8088 and SSH remain among the top exploited TCP ports for 2021. These are typically abused by IoT botnets, including many of the Mirai variants, that are continuing to wreak havoc on the internet through DDoS attacks and put IoT devices such as IP cameras and network devices such as routers and modems at risk. While Telnet was a Mirai favorite for a long time, the events on SSH surpassed Telnet by almost 18 times. Most SSH attacks consist of account takeover and Brute Force attempts. Leveraging default credentials or leaked credentials, attackers try to get unauthorized access to devices and systems and either move laterally across organizations’ networks, abuse the resources of cloud instances for cryptomining, leverage the foothold as jump host to anonymize targeted attacks or leverage the devices’ connectivity to perform DDoS attacks.

4. Remote Desktop Protocol (RDP) is a proprietary protocol developed by Microsoft that provides a user with a graphical interface to connect to another computer over a network connection. (Source: Wikipedia)

5. Virtual Network Computing (VNC) is a graphical desktop-sharing system that uses the Remote Frame Buffer (RFB) protocol to remotely control another computer. It transmits the keyboard and mouse input from one computer to another, relaying the graphical-screen updates, over a network. (Source: Wikipedia)

6. Server Message Block (SMB) is a communication protocol^[1] that Microsoft created for providing shared access to files and printers across nodes on a network. (Source: Wikipedia)

MOST-SCANNED AND MOST-ATTACKED UDP PORTS

SIP (port 5060) was the most-targeted UDP-based service in 2021. Port 5060 is used by many SIP-based VoIP phones and providers. VoIP remains critical to organizations to ensure their productivity; and, for this reason, it also makes the list of most-targeted services for DDoS attacks in 2021. Vulnerabilities and weak or default passwords in VoIP services allow them to be abused for initial access, spying and moving laterally inside organizations’ networks.

NTP (port 123), Memcached (port 11211), LDAP (port 389), SSDP/UPnP (port 1900), SNMP (port 161) and mDNS (port 5353) are among the most-leveraged protocols for DDoS amplification attacks. Many black- and white-hat actors are continuously scanning and cataloging the internet’s addressable range to abuse for DDoS attacks (black hat) or assess the risk in the DDoS threat landscape (white hat).

Microsoft SQL Server (port 1434) is used by the Microsoft SQL Server database management system monitor and abused through remote code execution vulnerabilities and known for the W32.Spybot.Worm that spread through SQL Server 2000 and SQL Server Desktop Engine 2000. It was still a very solicited port in 2021.

FIGURE 58:
Top scanned and attacked TCP ports






FIGURE 59:
Top scanned and attacked UDP ports



ORIGINATING COUNTRIES

The top countries in which unsolicited network activity originated in 2021 were the United States, Russia, China, the United Kingdom and the Netherlands. However, as mentioned earlier, the real origin of an attack can be spoofed to impersonate attacks from a different country.

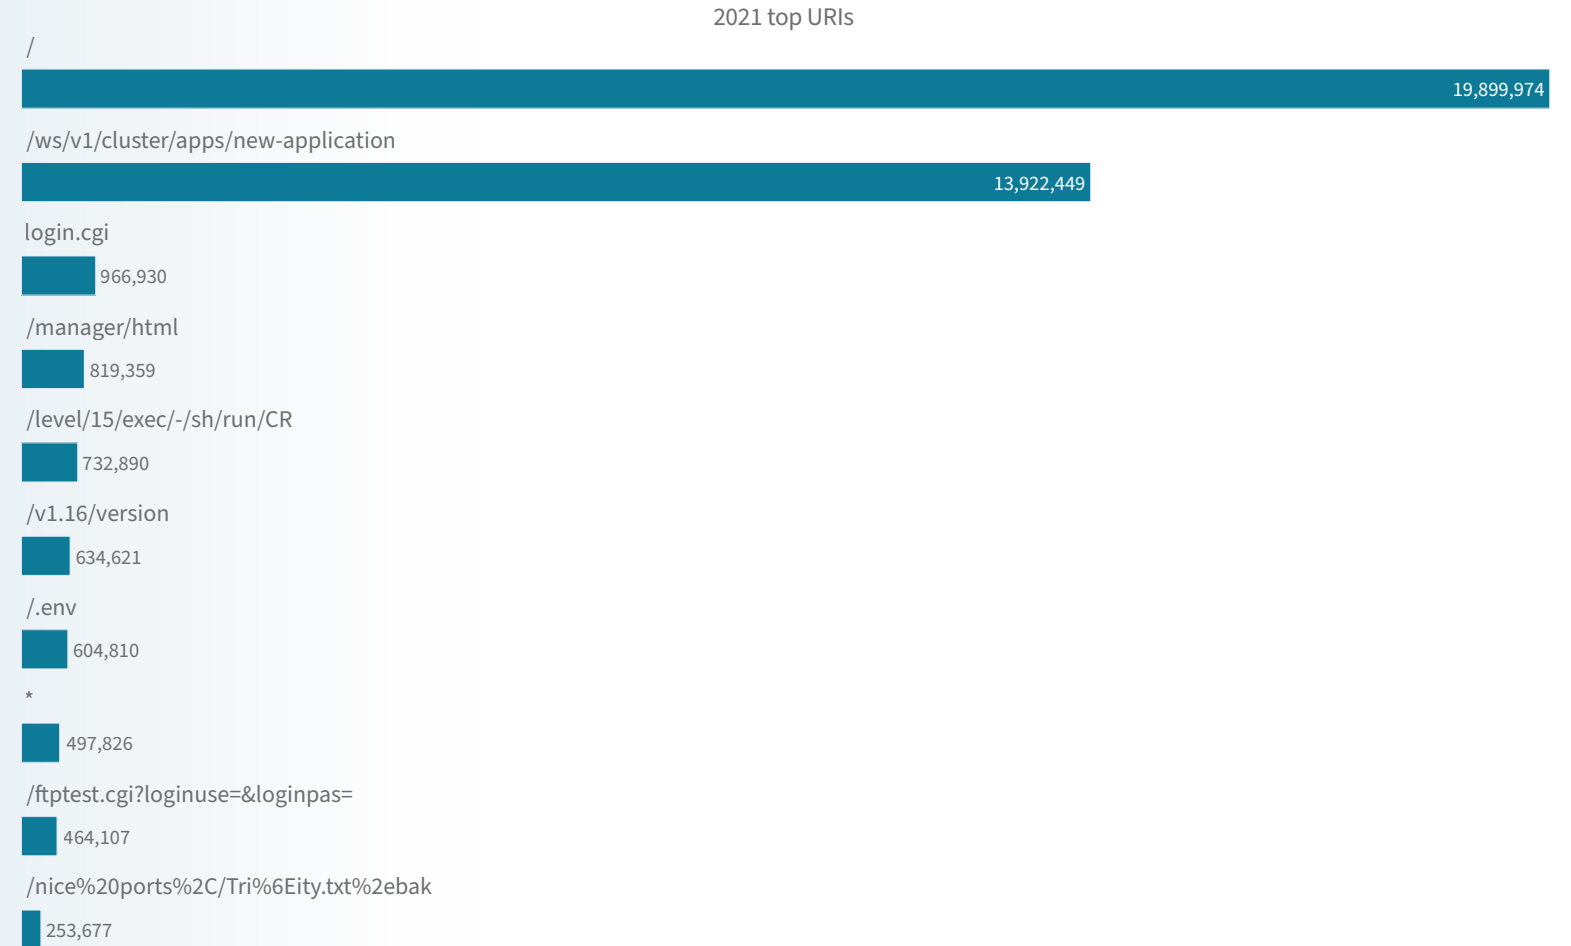
FIGURE 60:
Top attacking
countries



WEB SERVICE ATTACKS

The top attacked HTTP Uniform Resource Identifiers (URIs) are led by “/”, the universal URI for testing the presence of a web service and collecting information from header fields in server responses. There is a significant difference in the top targeted URIs for unsolicited events compared to top targets in web application attacks where services are backed by real applications. This section covers unsolicited events, meaning there is no real application or service running on the web server. The top URIs need to be interpreted as the top services and applications that are targeted by actors that are randomly scanning and exploiting the internet. Typically, a URI will conform to a known and disclosed vulnerability.

FIGURE 61:
Top scanned URIs



/ws/v1/cluster/app/new-application

A known vulnerability used to exploit Apache Hadoop YARN services and schedule arbitrary workloads on Hadoop clusters [71]. An exploit seen leveraged by many cryptojacking campaigns that try to leverage capable cloud instances of enterprises and research institutions illegitimately [72]. Was #1 exploit in 2020 [73].

/manager/html

Apache Tomcat Manager Application Upload Authenticated Code Execution vulnerability. This module can be used to execute a payload on Apache Tomcat servers that have an exposed “manager” application. The payload is uploaded as a WAR archive containing a JSP application using a POST request against the /manager/html/upload component. Was #2 exploit in 2020 [73].

/level/15/exec/-/sh/run/CR

Cisco routers have offered an HTTP interface since IOS release 11.2, first released in July 1999, that allows a user to execute commands directly from a URL. Attackers are still trying to find Cisco routers without authentication on the HTTP interface. Many routers have been deployed without changing default passwords or basic hardening practices, allowing for such opportunistic behavior by threat actors to bear fruit. Was #3 exploit in 2020 [73].

/v1.16/version

Used by threat actors to identify the available Docker API version through invoking a command for an old version. Used by cryptominers for abusing containers through the Docker API [74].

/ftptest.cgi?loginuse=&loginpas=

Known vulnerabilities in Wireless IP Camera (P2P) WIFICAM. Was #5 exploit in 2020 [73].

/nice%20ports%2C/Tri%6Eity.txt%2ebak

Request for “/nice ports,/Trinity.txt.bak” is used by Nmap’s service detection routine to test how a server handles escape characters within a URI.

TOP USER AGENTS


In HTTP, the User-Agent string is often used for content negotiation, where the origin server selects suitable content or operating parameters for the response. For example, the User-Agent string might be used by a web server to choose variants based on the known capabilities of a particular version of client software and differentiate its interface for different available screen real estate on mobile phones, tables and desktop browsers. The concept of content tailoring is built into the HTTP standard in RFC 1945 “for the sake of tailoring responses to avoid particular user agent limitations” [75].

As such, the User-Agent field in a web request can be used to identify the client agent that makes the request. Some malicious actors are aware of this identifying feature being leveraged to score the legitimacy of a web request by web security modules and mask their origins by randomly generating and changing the user agent to known legitimate values.

Commercial and open-source web-service-vulnerability scanning tools can be identified through their user agent, such as ZGrab, the application-layer network scanning component of the Zmap open-source scanning tool.

Some web crawlers and robots use the user agent to identify themselves. Websites can use a “robots.txt” file to regulate which search engine crawlers have access to which parts of the website. The “robots.txt” is a noncompulsory solution that relies completely on the crawler or robot. Needless to say, malicious bots will ignore the “robots.txt” entries and will crawl and scrape at their leisure. The Nimbostratus-Bot, for example, is considered a legitimate bot, and Cloud System Networks leverages the user agent to make its intentions clear by adding a URL to their homepage that explains the rationale behind its activity.

FIGURE 62:
Top User Agents




TOP HTTP CREDENTIALS

Not all web service vulnerabilities can be exploited without authenticating. Some web services have widely used default settings and some even have hard-coded secrets to protect access from unauthorized users or devices. The typical weak passwords combined in credential pairs with user admin or root were “admin”, “pass”, “password”, “123456”, “1234”, “1111”, “1234” and no password. These weak password permutations make up 8 of the top 10 credentials. These are universally agreed upon as the worst credentials and also the most abused because they provide a good amount of access to unauthorized devices that did not have their default credentials changed on installation.

“root:icatch99” is a hard-coded credential in DVRs from vendor LILIN that was publicly disclosed in March 2020 [1]. DVRs are ubiquitous in the IoT landscape, as are the security cameras that feed them.

FIGURE 63:
Top HTTP credentials

2021 top HTTP credentials




The credentials “8hYTSUFk:8hYTSUFk” are still a bit of a mystery. The base64 encoded version of the credential pair is “OGhZVFNVrms6OGhZVFNVrms=”. This authentication string was used in an example for passing authentication arguments to a generic web API interaction and exploration module for Node called Yiff Rewrite [2], an extended wrapper based on Hokkqi’s furry API wrapper, published in npm, the package manager for Node. The string was also discovered in several malware binaries, one of which being the Windows x86 (64-bit) console executable “l5obas.exe” [76].

TOP SSH USERNAMES

The top usernames used during SSH authentication provide information on the most-sought-for and most-likely services vulnerable to Brute Force. Among the top 10 are “postgres”, “oracle” and “git”.

FIGURE 64:
Top SSH usernames



References

- [1] C. Cimpanu, “DDoS Botnets Have Abused Three Zero-Days in LILIN Video Recorders for Months,” ZDNet, March 21, 2020. [Online]. Available: www.zdnet.com/article/ddos-botnets-have-abused-three-zero-days-in-lilin-video-recorders-for-months.
- [2] MrGriefs, “Yiff Rewrite, an Extended Wrapper Based on Hokkqi’s Furry API Wrapper,” Github, April 6, 2021. [Online]. Available: <https://github.com/MrGriefs/furry-wrapper>. [Accessed February 10, 2022].
- [3] “SolarWinds Orion Supply Chain Attack,” Radware, December 15, 2020. [Online]. Available: www.radware.com/security/ddos-threats-attacks/threat-advisories-attack-reports/solarwinds-orion-supply-chain-attack.
- [4] D. E. Sanger, C. Krauss and N. Perlroth, “Cyberattack Forces a Shutdown of a Top U.S. Pipeline,” The New York Times, May 8, 2021. [Online]. Available: www.nytimes.com/2021/05/08/us/politics/cyberattack-colonial-pipeline.html.
- [5] “World’s Most Dangerous Malware EMOTET Disrupted Through Global Action,” Europol, January 27, 2021. [Online]. Available: www.europol.europa.eu/media-press/newsroom/news/world%e2%80%99s-most-dangerous-malware-emotet-disrupted-through-global-action.
- [6] R. B. Yizhak, “The Re-Emergence of Emotet,” Deep Instinct, November 30, 2021. [Online]. Available: www.deepinstinct.com/blog/the-re-emergence-of-emotet.
- [7] “СБУ заблокувала діяльність транснаціонального хакерського угруповання (SBU Blocked the Activities of a Transnational Hacker Group),” SBU, February 17, 2021. [Online]. Available: <https://ssu.gov.ua/novyny/sbu-zablokuvala-diiialnist-transnatsionalnoho-khakerskoho-uhrupovannia>.
- [8] “ProxyLogon: Zero-Day Exploits in Microsoft Exchange Server,” Radware, March 16, 2021. [Online]. Available: www.radware.com/security/ddos-threats-attacks/threat-advisories-attack-reports/proxy-logon.
- [9] “HAFNIUM Targeting Exchange Servers with 0-Day Exploits,” Microsoft Threat Intelligence Center (MSTIC), March 2, 2021. [Online]. Available: www.microsoft.com/security/blog/2021/03/02/hafnium-targeting-exchange-servers.
- [10] M. Faou, M. Tartare and T. Dupuy, “Exchange Servers Under Siege from at Least 10 APT Groups,” Welivesecurity by ESET, March 10, 2021. [Online]. Available: www.welivesecurity.com/2021/03/10/exchange-servers-under-siege-10-apt-groups.
- [11] D. Palmer, “Microsoft Exchange Server Attacks: ‘They’re Being Hacked Faster Than We Can Count’, Says Security Company,” ZDNet, March 22, 2021. [Online]. Available: www.zdnet.com/article/microsoft-exchange-server-attacks-theyre-being-hacked-faster-than-we-can-count-says-security-company.
- [12] “F5 Security Advisory for RCE Vulnerabilities in BIG-IP, BIG-IQ,” CISA, March 10, 2021. [Online]. Available: www.cisa.gov/uscert/ncas/current-activity/2021/03/10/f5-security-advisory-rce-vulnerabilities-big-ip-big-iq.
- [13] L. O’Donnell, “Critical F5 BIG-IP Flaw Now Under Active Attack,” Threatpost, March 19, 2021. [Online]. Available: <https://threatpost.com/critical-f5-big-ip-flaw-now-under-active-attack/164940>.
- [14] C. Cimpanu, “Despite Arrests in Spain, FluBot Operations Explode Across Europe and Japan,” The Record, April 26, 2021. [Online]. Available: <https://therecord.media/despite-arrests-in-spain-flubot-operations-explode-across-europe-and-japan>.

- [15] “SA44784 - 2021-04: Out-of-Cycle Advisory: Multiple Vulnerabilities Resolved in Pulse Connect Secure 9.1R11.4,” Pulse Secure, April 2021. [Online]. Available: https://kb.pulsesecure.net/articles/Pulse_Security_Advisories/SA44784.
- [16] D. Perez, S. Jones, G. Wood and S. Eckels, “Check Your Pulse: Suspected APT Actors Leverage Authentication Bypass Techniques and Pulse Secure Zero-Day,” Mandiant, April 20, 2021. [Online]. Available: www.mandiant.com/resources/suspected-apt-actors-leverage-bypass-techniques-pulse-secure-zero-day.
- [17] A. Sharma, “Linux Bans University of Minnesota for Committing Malicious Code,” BleepingComputer, April 21, 2021. [Online]. Available: www.bleepingcomputer.com/news/security/linux-bans-university-of-minnesota-for-committing-malicious-code.
- [18] “Ransom DDoS Update: The Hunt for Unprotected Assets,” Radware, June 11, 2021. [Online]. Available: www.radware.com/security/ddos-threats-attacks/threat-advisories-attack-reports/ransom-ddos-update-hunt-for-unprotected-assets.
- [19] C. Cimpanu, “North Korean Hackers Breach South Korea’s Atomic Research Agency Through VPN Bug,” The Record, June 19, 2021. [Online]. Available: <https://therecord.media/north-korean-hackers-breach-south-koreas-atomic-research-agency-through-vpn-bug>.
- [20] “VPN Attacks Up Nearly 2000% as Companies Embrace a Hybrid Workplace,” Help Net Security, June 15, 2021. [Online]. Available: www.helpnetsecurity.com/2021/06/15/vpn-attacks-up.
- [21] “Anatomy of a Compromised Account,” Agari by HelpSystems, June 8, 2021. [Online]. Available: www.agari.com/insights/whitepapers/anatomy-compromised-account.
- [22] D. Palmer, “This Is How Fast a Password Leaked on the Web Will Be Tested Out by Hackers,” ZDNet, June 8, 2021. [Online]. Available: www.zdnet.com/article/this-is-how-fast-a-password-leaked-on-the-web-will-be-tested-out-by-hackers.
- [23] J. Cox, “How Hackers Used Slack to Break into EA Games,” Vice, June 11, 2021. [Online]. Available: www.vice.com/en/article/7kvkqb/how-ea-games-was-hacked-slack.
- [24] “Mass Scanning for VMWare vCenter RCE,” Radware, June 7, 2021. [Online]. Available: www.radware.com/security/ddos-threats-attacks/threat-advisories-attack-reports/mass-scanning-vmware-vcenter-rce.
- [25] A. Sharma, “Sonatype Catches New PyPI Cryptomining Malware,” Sonatype, June 21, 2021. [Online]. Available: <https://blog.sonatype.com/sonatype-catches-new-pypi-cryptomining-malware-via-automated-detection>.
- [26] “Kaseya VSA Ransomware Attack,” Wikipedia, 2021. [Online]. Available: https://en.wikipedia.org/wiki/Kaseya_VSA_ransomware_attack.
- [27] 360 Netlab, Twitter, July 28, 2021. [Online]. Available: <https://twitter.com/360Netlab/status/1420390398825058313>.
- [28] O. Yoachimik, “Cloudflare Thwarts 17.2M RPS DDoS Attack — the Largest Ever Reported,” Cloudflare, August 19, 2021. [Online]. Available: <https://blog.cloudflare.com/cloudflare-thwarts-17-2m-rps-ddos-attack-the-largest-ever-reported>.
- [29] K. Bock, A. Alaraj, Y. Fax, K. Hurley, E. Wustrow and D. Levin, “Weaponizing Middleboxes for TCP Reflected Amplification,” August 12, 2021. [Online]. Available: <https://geneva.cs.umd.edu/posts/usenix21-weaponizing-censors>.

- [30] GetNews, “IPStress Offers One of the Finest DDoS for Hire Service,” Digital Journal, August 10, 2021. [Online]. Available: www.digitaljournal.com/pr/ipstress-offers-one-of-the-finest-ddos-for-hire-service.
- [31] “Dark.IoT Botnet,” Radware, August 24, 2021. [Online]. Available: www.radware.com/security/threat-advisories-and-attack-reports/dark-iot-botnet.
- [32] “By Design: How Default Permissions on Microsoft Power Apps Exposed Millions,” UpGuard, August 23, 2021. [Online]. Available: www.upguard.com/breaches/power-apps.
- [33] C. Cimpanu, “1.9 Million Records from the FBI’s Terrorist Watchlist Leaked Online,” The Record, August 16, 2021. [Online]. Available: <https://therecord.media/1-9-million-records-from-the-fbis-terroris-watchlist-leaked-online>.
- [34] L. Abrams, “REvil Ransomware Is Back in Full Attack Mode and Leaking Data,” BleepingComputer, September 11, 2021. [Online]. Available: www.bleepingcomputer.com/news/security/revil-ransomware-is-back-in-full-attack-mode-and-leaking-data.
- [35] T. Richardson, “UK VoIP Telco Receives ‘Colossal Ransom Demand’, Reveals Revil Cybercrooks Suspected of ‘Organised’ DDoS Attacks on UK Voip Companies,” The Register, September 2, 2021. [Online]. Available: www.theregister.com/2021/09/02/uk_voip_telcos_revil_ransom.
- [36] “Mēris Botnet, Climbing to the Record,” Qrator Labs, September 9, 2021. [Online]. Available: https://blog.qrator.net/en/meris-botnet-climbing-to-the-record_142.
- [37] N. Ohfeld, “OMIGOD: Critical Vulnerabilities in OMI Affecting Countless Azure Customers,” WIZ, September 14, 2021. [Online]. Available: <https://blog.wiz.io/omigod-critical-vulnerabilities-in-omi-azure>.
- [38] “Dark.IoT, OMIGOD & UDP Technology Update,” Radware, September 21, 2021. [Online]. Available: www.radware.com/security/threat-advisories-and-attack-reports/dark-iot-omigod-update.
- [39] I. A. Titouan Lazard, “UDP Technology IP Camera Vulnerabilities,” Randorisec, July 8, 2021. [Online]. Available: www.randorisec.fr/udp-technology-ip-camera-vulnerabilities.
- [40] “Illinois Man Convicted of Federal Criminal Charges for Operating Subscription-Based Computer Attack Platforms,” The United States Attorney’s Office, Central District of California, September 16, 2021. [Online]. Available: www.justice.gov/usao-cdca/pr/illinois-man-convicted-federal-criminal-charges-operating-subscription-based-computer.
- [41] “VMware vCenter Server Updates Address Multiple Security Vulnerabilities,” VMWare, September 21, 2021. [Online]. Available: www.vmware.com/security/advisories/VMsa-2021-0020.html.
- [42] C. Duckett, “RCE Is Back: VMware Details File Upload Vulnerability in vCenter Server,” ZDNet, 22 September 2021. [Online]. Available: www.zdnet.com/article/rce-is-back-vmware-details-file-upload-vulnerability-in-vcenter-server.
- [43] J. Greig, “DDoS Attack Cost Bandwidth.com Nearly \$12 Million,” ZDNet, November 8, 2021. [Online]. Available: www.zdnet.com/article/ddos-attack-cost-bandwidth-com-nearly-12-million.
- [44] P. Nederland, “Kopers van DDoS-Aanval Krijgen Waarschuwing van Cybercrimeteam,” Politie.nl, October 11, 2021. [Online]. Available: www.politie.nl/nieuws/2021/oktober/11/03-kopers-van-ddos-aanval-krijgen-waarschuwing-van-cybercrimeteam.html.

- [45] C. Cimpanu, “REvil Gang Shuts Down for the Second Time After Its Tor Servers Were Hacked,” The Record, October 18, 2021. [Online]. Available: <https://therecord.media/revil-gang-shuts-down-for-the-second-time-after-its-tor-servers-were-hacked>.
- [46] J. Menn and C. Bing, “Governments Turn Tables on Ransomware Gang REvil by Pushing It Offline,” Reuters, October 22, 2021. [Online]. Available: www.reuters.com/technology/exclusive-governments-turn-tables-ransomware-gang-revil-by-pushing-it-offline-2021-10-21.
- [47] “Threat Advisory: Apache HTTP Server Zero-Day Vulnerability Opens Door for Attackers,” Cisco Talos, October 7, 2021. [Online]. Available: <https://blog.talosintelligence.com/2021/10/apache-vuln-threat-advisory.html>.
- [48] bashis, “Dahua Authentication Bypass.txt,” Github, October 6, 2021. [Online]. Available: <https://github.com/mcw0/PoC/blob/master/Dahua%20authentication%20bypass.txt>.
- [49] IPVM Team, “Best Buy, Home Depot and Lowes Drop Dahua Lorex,” IPVM, October 25, 2021. [Online]. Available: <https://ipvm.com/reports/lorex-box>.
- [50] J. Honovich and C. Rollet, “US FCC: Dahua and Hikvision ‘Deemed a Threat’ to National Security,” IPVM, March 15, 2021. [Online]. Available: <https://ipvm.com/reports/fcc-hikua>.
- [51] J. Honovich, “Hikvision and Dahua Sanctioned for Human Rights Abuses,” IPVM, October 7, 2019. [Online]. Available: <https://ipvm.com/reports/sanction-hikua>.
- [52] E. Elkin and D. Shanker, “That Cream Cheese Shortage You Heard About? Cyberattacks Played a Part,” Bloomberg, December 9, 2021. [Online]. Available: www.bloomberg.com/news/articles/2021-12-09/that-cream-cheese-shortage-you-heard-about-cyberattacks-played-a-part.
- [53] T. Starks, “‘Cyber Event’ Knocks Dairy Giant Schreiber Foods Offline amid Industry Ransomware Outbreak,” Cyberscoop, October 27, 2021. [Online]. Available: www.cyberscoop.com/schreiber-foods-cyber-event-ransomware-agriculture-food.
- [54] FBI, “Cyber Criminal Actors Targeting the Food and Agriculture Sector with Ransomware Attacks,” DHS-CISA, September 1, 2021. [Online]. Available: <https://s3.documentcloud.org/documents/21053966/fbi-bc-cyber-criminal-actors-targeting-the-food-and-agriculture-sector-with-ransomware-attacks.pdf>.
- [55] C. Cimpanu, “TP-Link Routers Under Attack from Dark.IoT Botnet,” The Record, December 9, 2021. [Online]. Available: <https://therecord.media/tp-link-routers-under-attack-from-dark-iot-botnet>.
- [56] “TCP Reflection Attacks,” Radware, November 7, 2019. [Online]. Available: www.radware.com/security/ddos-threats-attacks/threat-advisories-attack-reports/tcp-reflection-attacks.
- [57] “Memcached Under Attack,” Radware, March 1, 2018. [Online]. Available: www.radware.com/security/ddos-threats-attacks/threat-advisories-attack-reports/memcached-under-attack.
- [58] A. Toh, “Azure DDoS Protection—2021 Q3 and Q4 DDoS Attack Trends,” Microsoft, January 25, 2022. [Online]. Available: <https://azure.microsoft.com/en-us/blog/azure-ddos-protection-2021-q3-and-q4-ddos-attack-trends>.
- [59] T. Emmons, “Largest Ever Recorded Packet per Second-Based DDoS Attack Mitigated by Akamai,” Akamai, June 25, 2020. [Online]. Available: www.akamai.com/blog/news/largest-ever-recorded-packet-per-secondbased-ddos-attack-mitigated-by-akamai.
- [60] A. Dahan, “Business as Usual for Azure Customers Despite 2.4 Tbps DDoS Attack,” Microsoft, October 11, 2021. [Online]. Available: <https://azure.microsoft.com/en-us/blog/business-as-usual-for-azure-customers-despite-24-tbps-ddos-attack>.

- [61] A. Toh, “Azure DDoS Protection—2021 Q3 and Q4 DDoS Attack Trends,” Microsoft, January 25, 2022. [Online]. Available: <https://azure.microsoft.com/en-us/blog/azure-ddos-protection-2021-q3-and-q4-ddos-attack-trends>.
- [62] “Ransom DDoS Campaign: Circling Back,” Radware, January 22, 2021. [Online]. Available: www.radware.com/security/ddos-threats-attacks/threat-advisories-attack-reports/ddos-extortions-back.
- [63] D. Smith, “Welcome to the New World of Triple Extortion Ransomware,” Security Magazine, May 18, 2021. [Online]. Available: www.securitymagazine.com/articles/95238-welcome-to-the-new-world-of-triple-extortion-ransomware.
- [64] “Threat Advisory: Apache HTTP Server Zero-Day Vulnerability Opens Door for Attackers,” Cisco Talos, October 7, 2021. [Online]. Available: <https://blog.talosintelligence.com/2021/10/apache-vuln-threat-advisory.html>.
- [65] “Log4Shell: Critical Log4j Vulnerability,” Radware, December 12, 2021. [Online]. Available: www.radware.com/security/threat-advisories-and-attack-reports/log4shell-critical-log4j-vulnerability.
- [66] “Mozi, Another Botnet Using DHT,” 360 Netlab, December 23, 2019. [Online]. Available: <https://blog.netlab.360.com/mozi-another-botnet-using-dht>.
- [67] P. Geenens, “Hajime – Sophisticated, Flexible, Thoughtfully Designed and Future-Proof,” Radware, April 26, 2017. [Online]. Available: <https://blog.radware.com/security/2017/04/hajime-futureproof-botnet>.
- [68] O. Harpaz, “FritzFrog: A New Generation of Peer-to-Peer Botnets,” Guardicore, 2020. [Online]. Available: www.guardicore.com/labs/fritzfrog-a-new-generation-of-peer-to-peer-botnets.
- [69] V. Singhal, R. Nigam, Z. Zhang and A. Davila, “New Mirai Variant Targeting Network Security Devices,” Palo Alto Networks Unit 42, March 15, 2021. [Online]. Available: <https://unit42.paloaltonetworks.com/mirai-variant-iot-vulnerabilities>.
- [70] D. Fiser, “More Than 8,000 Unsecured Redis Instances Found in the Cloud,” Trend Micro, April 2, 2020. [Online]. Available: www.trendmicro.com/en_us/research/20/d/more-than-8-000-unsecured-redis-instances-found-in-the-cloud.html.
- [71] P. Geenens, “Hadoop YARN: An Assessment of the Attack Surface and Its Exploits,” Radware, November 15, 2018. [Online]. Available: <https://blog.radware.com/security/2018/11/hadoop-yarn-an-assessment-of-the-attack-surface-and-its-exploits>.
- [72] “Demonbot,” Radware, October 25, 2018. [Online]. Available: www.radware.com/security/ddos-threats-attacks/threat-advisories-attack-reports/demonbot.
- [73] Radware Vulnerability Research Team, “The Top Web Service Exploits in 2020,” Radware, December 23, 2020. [Online]. Available: <https://blog.radware.com/security/2020/12/the-top-web-service-exploits-in-2020>.
- [74] Y. Chikvashvili, “Cryptocurrency Miners Abusing Containers: Anatomy of an (Attempted) Attack,” Aqua, February 15, 2018. [Online]. Available: <https://blog.aquasec.com/cryptocurrency-miners-abusing-containers-anatomy-of-an-attempted-attack>.
- [75] “User agent,” Wikipedia, September 30, 2021. [Online]. Available: https://en.wikipedia.org/wiki/User_agent.
- [76] “Static and Dynamic Analysis of Malicious Binary I5obas.exe,” Hybrid Analysis, September 18, 2021. [Online]. Available: www.hybrid-analysis.com/sample/1b37c8074dbfad28b7dc8c1dc631d3f16cbcd58b93aa308386a69159aadabb99/615357ca878ee204921a55ef. [Accessed February 10, 2022].

List of Figures and Tables

| | | | |
|---|----|--|----|
| Figure 1: Blocked malicious events, normalized per customer | 13 | Figure 25: Number of attacks larger than 10Gbps, normalized per 1,000 attacks | 23 |
| Figure 2: Blocked volume, normalized per customer | 13 | Figure 26: Number of attacks larger than 100Gbps, normalized per 1,000 attacks ... | 23 |
| Figure 3: Blocked malicious events, normalized per customer | 13 | Figure 27: Top protocols leveraged by attacks in 2021 (by packets)..... | 24 |
| Figure 4: Blocked volume, normalized per customer | 13 | Figure 28: Top attack vectors in 2021 (by packets)..... | 25 |
| Figure 5: Number of DDoS attacks mitigated per quarter in 2021 | 14 | Figure 29: Top attacked application protocols in 2021 (by packets)..... | 25 |
| Figure 6: Cumulative sum of DDoS attacks per day throughout 2021 | 14 | Figure 30: Amplification vectors per quarter by volume | 25 |
| Figure 7: Quarter-to-quarter change in the number of attacks per customer..... | 14 | Figure 31: Linear relationship between attack vector throughput and rate in function of packet size..... | 26 |
| Figure 8: Average DDoS attack volume, normalized per customer | 15 | Figure 32: Global diversity in attack vectors in function of attack vector size..... | 28 |
| Figure 9: Average volume per DDoS attack, normalized..... | 15 | Figure 33: Average packet size in function of attack vector size | 29 |
| Figure 10: Average and maximum attack sizes..... | 15 | Figure 34: Average vector duration per attack vector size..... | 29 |
| Figure 11: Blocked volume per region, normalized, for 2020 and 2021..... | 16 | Figure 35: Average volume by attack vector size | 30 |
| Figure 12: Blocked volume per industry, normalized, for 2020 and 2021 | 16 | Figure 36: Average volume per attack vector, by vector size, log scale | 31 |
| Figure 13: Yearly volume per industry, normalized | 17 | Figure 37: Average packet rate per attack vector, by vector size, log scale..... | 32 |
| Figure 14: Yearly DoS events per industry, normalized..... | 18 | Figure 38: Number of attacks per attack vector size | 33 |
| Figure 15: Volume per event by industry..... | 18 | Figure 39: Relative share of UDP and TCP attack vectors | 33 |
| Figure 16: Quarterly DoS events per industry, normalized..... | 19 | Figure 40: Total attack volume per vector size | 33 |
| Figure 17: Quarterly blocked volume per industry, normalized | 19 | Figure 41: Number of distinct attack vectors per attack in function of attack size..... | 34 |
| Figure 18: Quarterly number of large attacks | 20 | Figure 42: Average attack duration in function of attack size | 34 |
| Figure 19: Yearly number of attack vectors larger than 10Gbps | 20 | Figure 43: Average volume per attack in function of attack size..... | 35 |
| Figure 20: Quarterly number of mid-sized attacks..... | 21 | Figure 44: Malicious events by attack category | 38 |
| Figure 21: Yearly number of mid-sized attacks..... | 22 | Figure 45: Top network intrusions in 2020 versus 2021 | 39 |
| Figure 22: Quarterly number of micro flood attacks..... | 22 | Figure 46: Daily blocked Log4Shell activity in Radware’s Cloud WAF Service and Cloud DDoS Protection Services..... | 42 |
| Figure 23: Yearly number of micro flood attacks..... | 22 | | |
| Figure 24: Number of attacks larger than 1Gbps, normalized per 1,000 attacks | 23 | | |

Figure 47: Yearly blocked web application transactions 43

Figure 48: Quarterly blocked web application transactions 43

Figure 49: 2021 web application blocked transactions versus attacks
blocked by signature 43

Figure 50: Top security violation types, normalized per customer..... 44

Figure 51: Violation types for known web application attacks by quarter 44

Figure 52: Blocked security violations by 2017 OWASP Top 10 application
security risks 45

Figure 53: Web application attacks per 2017 OWASP Top 10 category..... 45

Figure 54: Top attacking countries..... 46

Figure 55: Web application attacks by industry 46

Figure 56: Number of events per month, recorded by
Radware’s Global Deception Network..... 49

Figure 57: Number of unique IPs per month, registered by
Radware’s Global Deception Network..... 49

Figure 58: Top scanned and attacked TCP ports..... 51

Figure 59: Top scanned and attacked UDP ports 51

Figure 60: Top attacking countries..... 52

Figure 61: Top scanned URIs..... 53

Figure 62: Top User Agents 55

Figure 63: Top HTTP credentials 56

Figure 64: Top SSH usernames..... 56

Table 1: Change in volume per customer per industry from 2020 to 2021 17

Table 2: DDoS amplification attack vectors..... 24

Table 3: Average volume per attack vector..... 30

Table 4: Average packet rate per attack vector 31

Methodology and Sources

The data for DDoS events and volumes was collected from a sampled set of Radware devices deployed in Radware cloud scrubbing centers and on-premise managed devices in Radware hybrid and peak protection services.

Radware’s Global Deception Network provides detailed events and payload data on a wide range of attacks and serves as a basis for the “Unsolicited Network Scanning and Attack Activity” section.

The data for web application attacks was collected from blocked application security events from the Radware Cloud WAF Service. Collected events were based solely on automatically detected and known vulnerability exploits and exclude any events that might be blocked or reported by custom rules added to a web application policy by managed services and/or customers.

About Radware

Radware® (NASDAQ: RDWR) is a global leader of cybersecurity and application delivery solutions for physical, cloud and software-defined data centers. Its award-winning solutions portfolio secures the digital experience by providing infrastructure, application and corporate IT protection and availability services to enterprises globally. Radware’s solutions empower more than 12,500 enterprise and carrier customers worldwide to adapt quickly to market challenges, maintain business continuity and achieve maximum productivity while keeping costs down. For more information, please visit www.radware.com.

Radware encourages you to join our community and follow us on: Radware Blog, LinkedIn, Facebook, Twitter, SlideShare, YouTube, Radware Connect app for iPhone® and our Security Research Center that provides a comprehensive analysis of DDoS attack tools, trends and threats. This document is provided for information purposes only.

This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law. Radware specifically disclaims any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. The technologies, functionalities, services or processes described herein are subject to change without notice.

Editors

Pascal Geenens – Director of Threat Intelligence
Daniel Smith – Head of Threat Research

Executive Sponsors

Shira Sagiv – VP Portfolio Marketing
Ron Meyran – Sr Director of Corporate Enablement

Production

Colin Beasty – Corporate Marketing Manager
Dasnet Garcia – Brand Marketing Manager
Gerri Dyrek – Director of Public Relations