radware
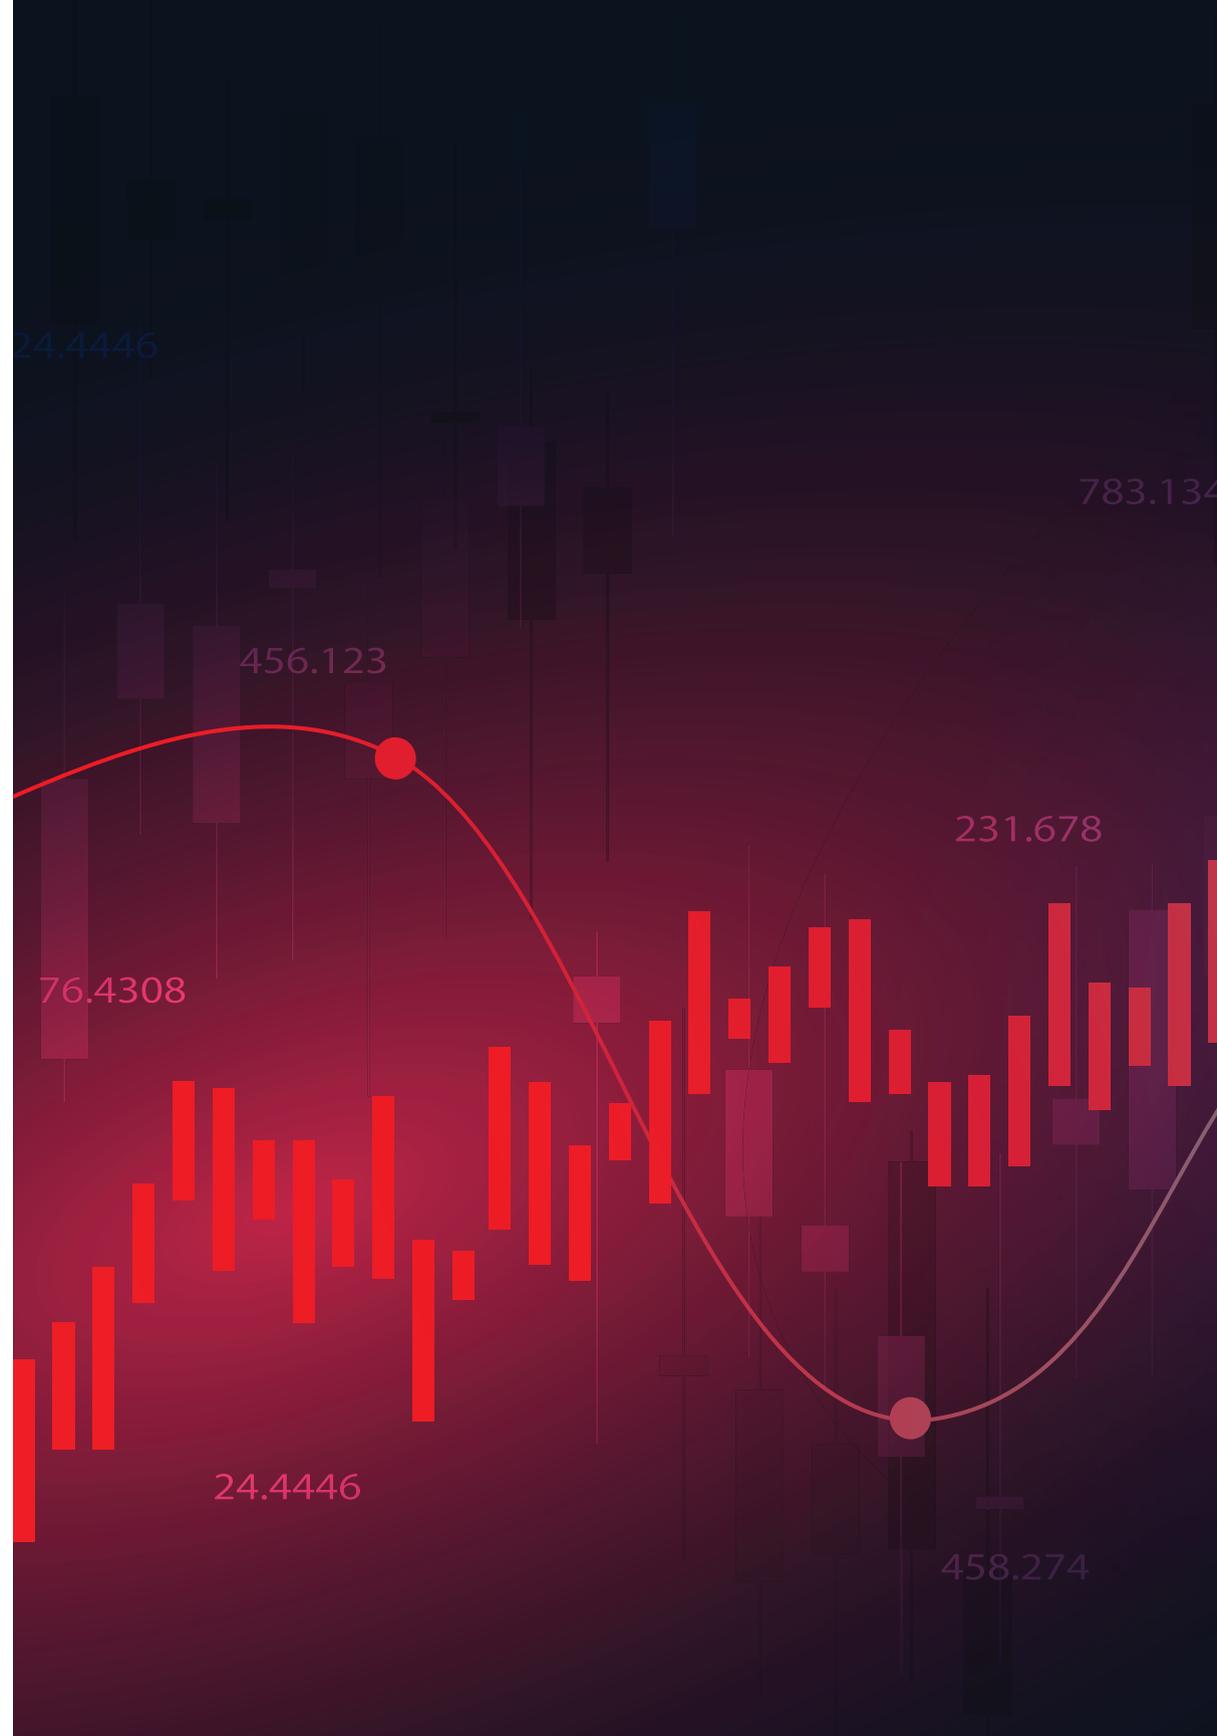
Quarterly DDoS and Application Attack Report

Radware's Quarterly DDoS and Application Attack Report provides an overview of attack activity witnessed during the third quarter of the 2021 calendar year. It analyzes network and application attack activity sourced from Radware's Cloud and Managed Services and Radware's Global Deception Network.

Contents

Executive Summary

DDOS ATTACK ACTIVITY

Fewer attacks during the summer months and a decrease in the average size of the attacks resulted in the lowest recorded total blocked volume in the last two years. The total number of events for the quarter was slightly below previous quarters in 2021 while staying above the level of the highest quarter in 2020. Notwithstanding, the total volume blocked in the first three quarters of 2021 increased by 44% compared to the same period in 2020 and is just 6% shy of the 2020 total volume. By August 2021, the number of malicious events blocked already exceeded the totals of 2020. During the first nine months of 2021, 75% more events were blocked compared to the same period in 2020.

UDP Fragmentation attacks were still responsible for the bulk of the attack volume, but SSL-ClearText traffic now accounted for a larger portion of the volume compared to UDP Floods. In Q1 and Q2 of 2021, UDP-based attacks accounted for 99% and 97%, respectively, of the total blocked volume. In Q3 of 2021, UDP-based attacks accounted for 88.4% of the total attack volume and TCP-based attacks represented over 10% of that total volume.

The low total volume in Q3 compared to the other quarters in 2021 and the change in targeted applications, protocols and attack vectors illustrate a shift in DDoS attacker tactics from saturation-based floods to server resource–consuming, application-level attacks.

APAC witnessed the fewest malicious events in Q3 but accounted for the majority of the blocked volume. In contrast, EMEA was the most attacked region but had the lowest total blocked volume per customer.

The most attacked industries in Q3 of 2021 were technology, healthcare and communications. Gaming and telecom had to endure over 50% of the total volume in Q3. Technology, research and education as well as finance and healthcare accounted for most of the remaining volume. Research and education saw its largest volume in September, which is not atypical when schools get back in session.

A single quarter is not an indicator of a trend. Our customers were spared from large volumetric attacks, but there is still a steady underflow of smaller attack activity, which caused the number of attacks to decrease only slightly. While Radware customers were spared from large-scale assaults, our industry in Q3 was not without its share of record DDoS attacks. VoIP telcos and other DDoS mitigation organizations were less fortunate. August was a month where DDoS attack records were challenged and broken across three major continents, and September marked the return of education DDoS attacks and service provider DDoS attacks on VoIP telecommunication providers in the United Kingdom and Canada that came with colossal ransom demands by an actor posing as "REvil."

WEB APPLICATION ATTACK ACTIVITY

Web application attacks based on known vulnerabilities and techniques are ramping up quickly, doubling every quarter this year. The low-hanging fruit represented by predictable resource location and injection attacks are the most prominent security violations blocked by the Radware application security services. Cross-site scripting (XSS) and information leaks close out the top five most often blocked violations.

The most offenders in Q3 were located in the United States and Russia. India, the United Kingdom and Germany completed the top five for the quarter.

Banking and finance was hit the hardest and accounted for almost 23% of all blocked web application security events. Government (16%), technology (15%) and retail (12%) were among the most attacked industries.

The top violations reported in Q3 are aligned with the top web application security risks published by the OWASP Foundation in the 2017 and 2021 OWASP Top 10 lists.

UNSOLICITED NETWORK SCANNING AND ATTACK ACTIVITY

Q3 activity peaked at 27 million events per day and almost 300 million events per month in August.

The top attacking countries based on arbitrary client IP information during Q3 were the United States followed by Russia, China, the United Kingdom and the Netherlands. Considering only nonspoofed client IP addresses that can be unmistakably tied to the offending device and its geolocation, the top attacking countries were China followed by the United States, Brazil, Russia and India.

The most scanned and attacked TCP services are SSH followed by VNC and RDP. Telnet, HTTP and HTTPS on different ports remain among the top exploited TCP services in Q3. These are typically abused by IoT botnets, including many Mirai variants, that are continuing to wreak havoc on the internet through DDoS attacks and put IoT devices such as IP cameras and home routers and modems at risk. While Telnet was a Mirai favorite for a long time, SSH overshadowed Telnet by 15 times in Q3.

Redis, an open source in-memory data store used as a database, cache and message broker climbed the ranks to sixth after a remote command execution vulnerability was disclosed in July. In April 2020, it was reported that more than 8,000 unsecured Redis instances were deployed in public clouds [1].

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 crypto mining, leverage the foothold as jump host to anonymize targeted attacks or leverage device connectivity to perform DDoS attacks.

The most attacked UDP service was SIP, a service associated with internet voice services and applications such as VoIP phones and providers. Considering the ransom DoS attacks targeting VoIP providers in the second half of Q3, this activity might reflect a correlation with the discovery activity by the actors. Vulnerabilities in VoIP services are also candidates for initial access and moving laterally inside organizations' networks, ultimately falling victim to ransomware or backdoors.

The remainder of the most attacked UDP ports were related to scanning activity for services such as NTP, Memcached, LDAP, SNMP, SSDP and DNS. These services, when incorrectly configured, can be abused to perform volumetric DDoS amplification and reflection attacks.

The top web application exploit targeted by actors randomly scanning and exploiting services is the Apache Hadoop YARN exploit [2], disclosed in October 2018. This exploit is seen leveraged by many cryptojacking campaigns that try to use capable on-premise or cloud-hosted Hadoop clusters of enterprises and research institutions illegitimately [3].

Denial-of-Service Attack Activity
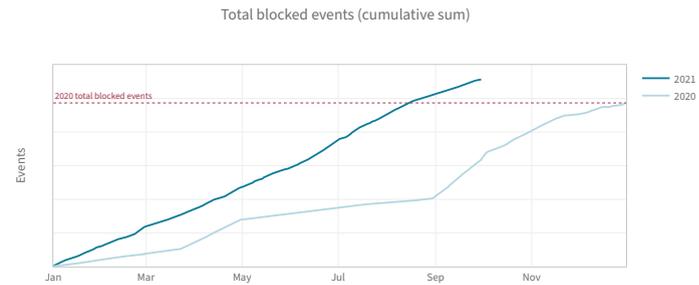
YEARLY TRENDS

By August 2021, the number of blocked malicious events exceeded the total number of malicious events blocked in 2020. During the first nine months of 2021, 75% more events were blocked compared to the same period in 2020.

The total volume blocked in the first three quarters of 2021 increased with 44% compared to the same period in 2020 and is just 6% below the total yearly volume blocked in 2020.

FIGURE 1:
Total number of blocked malicious events per year

Blocked malicious events
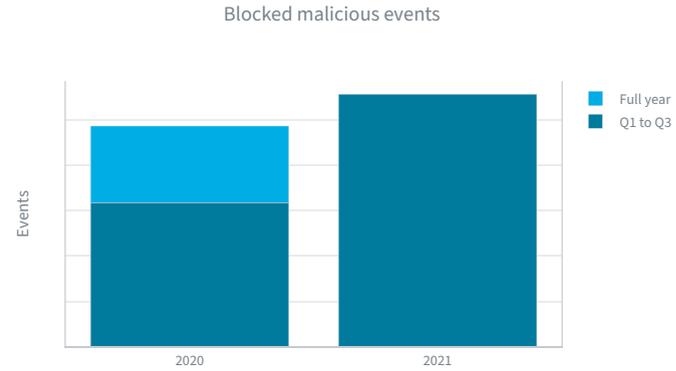
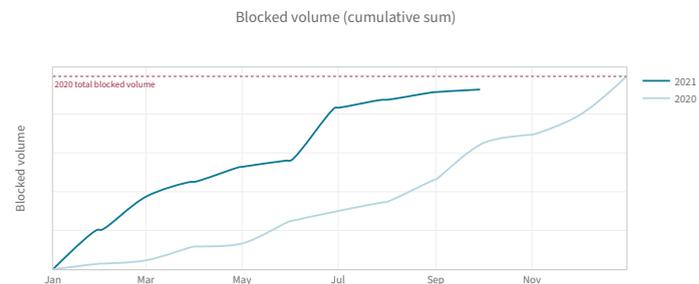**FIGURE 2:**
Total blocked events, cumulative sum over year

Total blocked events (cumulative sum)

FIGURE 3:
Total blocked volume per year
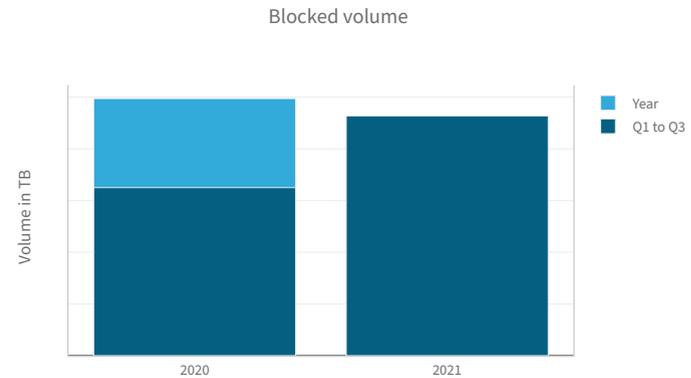
Blocked volume

FIGURE 4:
Blocked volume, cumulative sum over year

Blocked volume (cumulative sum)

QUARTERLY TRENDS

Compared to Q3 of 2020, the average number of blocked malicious events per customer increased by almost 20% to 12,400 events. The average blocked volume per customer dropped to 1.6TB per customer compared to 6.35TB in Q3 of 2020.

Compared to the first half of 2021, the average number of events per customer in Q3 decreased by 10%. The average blocked volume per customer, however, decreased by almost 80% from an average of 7.8TB per quarter in the first half to 1.6TB in Q3.

The average attack[1] size has decreased to 115Mbps after a significant growth in average attack size during the first and second quarter of 2021 and is slightly below the average attack size recorded in Q4 of 2020. The largest attack recorded in Q3 was 228Gbps, significantly lower than the 348Gbps attack recorded in Q2 of 2021.

1. Attacks are groups of one or more malicious events, overlapping in start time and duration, all representing a common, perceived attack on a customer. Attacks consist of one to hundreds of events, depending on the complexity and duration of the attack.

FIGURE 5:
Blocked malicious events, normalized per customer

Blocked malicious events (normalized)
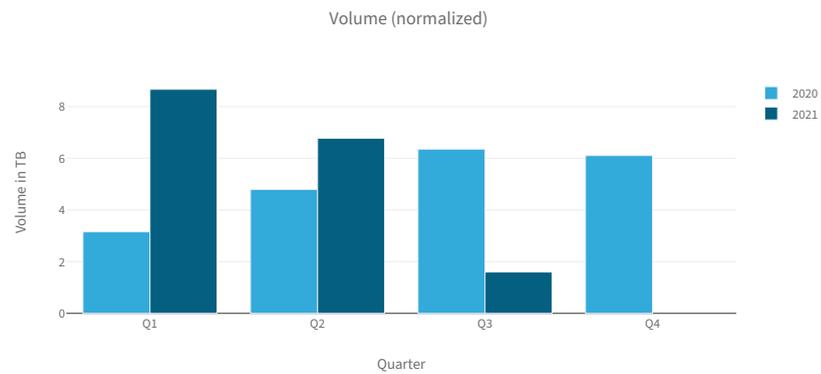
FIGURE 6:
Blocked malicious events, normalized per customer

Blocked malicious events (normalized)
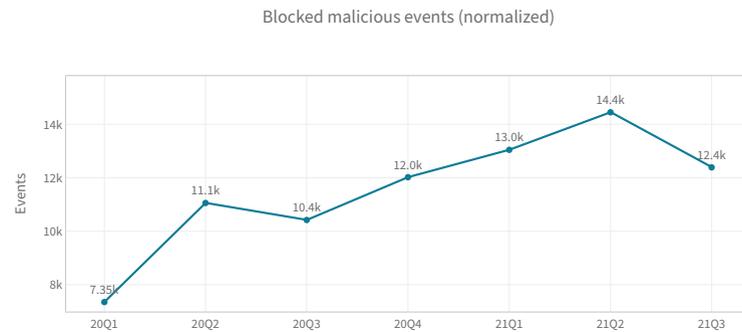
FIGURE 7:
Blocked volume in TB, normalized per customer

Volume (normalized)

FIGURE 8:
Blocked volume in TB, normalized per customer

Volume (normalized)

The average number of attacks per customer reached its year low during the summer months of July and August and is ramping back up in September to comparable levels with January, April and May of this year.

The relative number of attacks larger than 10Gbps is down from 3.31 per 1,000 attacks in Q2 to 1.87 per 1,000 attacks in Q3. The number of attacks larger than 1Gbps almost halved from 9.17 in Q2 to 4.72 per 1,000 attacks in Q3.

The decrease in average attack sizes correlates with the low volume of Q3, also taking into account that the number of malicious events only slightly dropped compared to other quarters in 2021.

FIGURE 9:
Average and maximum attack sizes

Attacks and attack sizes
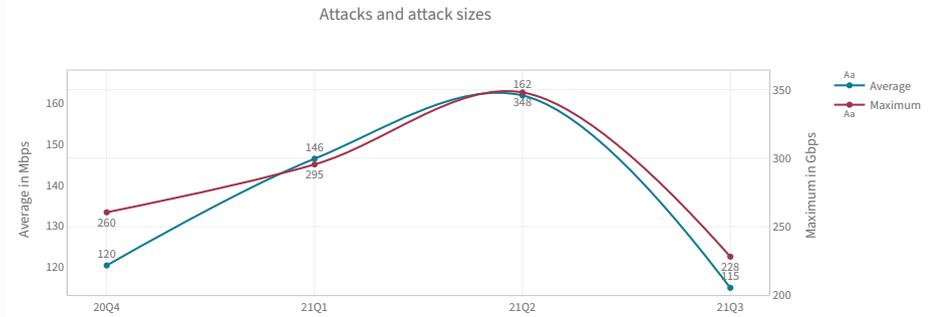
FIGURE 10:
Average number of attacks per customer

Average number of attacks per customer
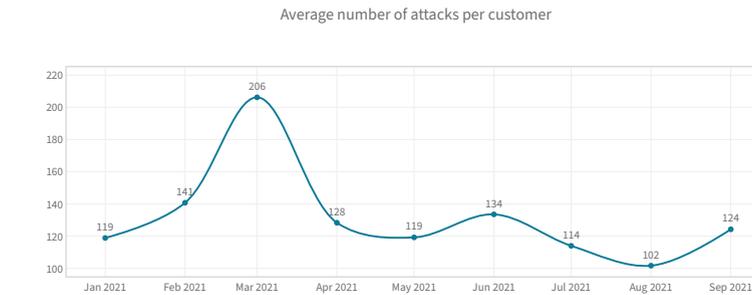
FIGURE 11:
Number of attacks larger than 10Gbps, normalized per 1,000 attacks

>10Gbps (normalized per 1,000 attacks)
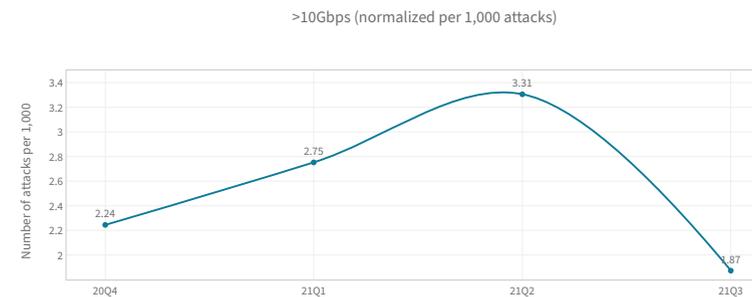
FIGURE 12:
Number of attacks larger than 1Gbps, normalized per 1,000 attacks

>10Gbps (normalized per 1,000 attacks)

REGIONS AND INDUSTRIES

The majority of blocked malicious events per customer was split equally between EMEA and the Americas in Q3, with less than 10% of the events being attributed to APAC. The majority of the blocked volume, normalized per customer, however, was consumed by our scrubbing centers protecting APAC customers. EMEA, while having a considerable amount of blocked events, accounted for the smallest part of the total blocked volume per customer.

The shift in volume toward APAC customers in Q3 redistributed the total blocked volume across the year more equally between the three regions.

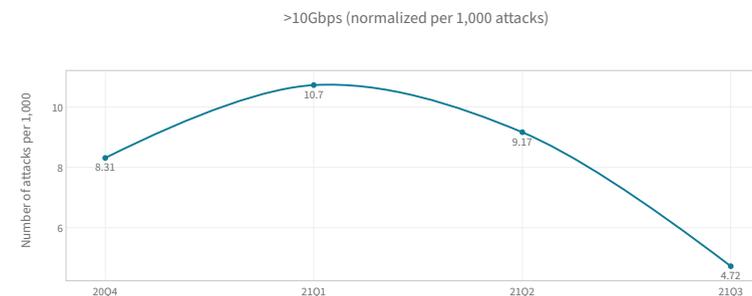The most attacked industry in Q3 was technology, with an average of 2,638 attacks per customer, followed by healthcare (1,785 attacks per customer), communications (1,525 attacks per customer), finance (1,337 attacks per customer), automotive (883 attacks per customer), gaming (598 attacks per customer), retail (556 attacks per customer), telecom (381 attacks per customer) and manufacturing (281 attacks per customer).

FIGURE 13:
Blocked events per region, normalized per customer

Blocked events (normalized)
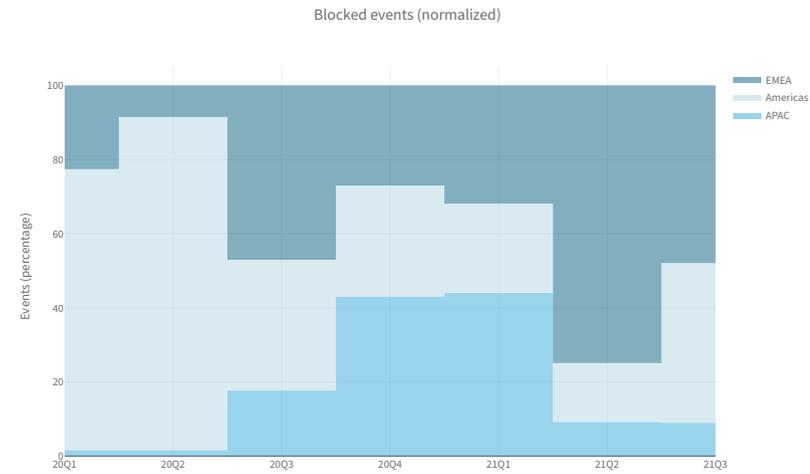
FIGURE 14:
Blocked volume per region, normalized per customer

Blocked volume (normalized)

FIGURE 15:
Blocked volume per region for 2021, normalized per customer
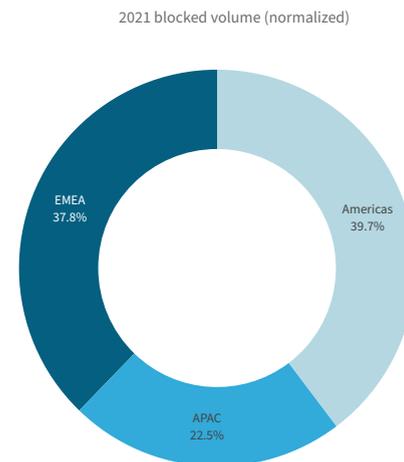
2021 blocked volume (normalized)

EMEA 37.8%
Americas 39.7%
APAC 22.5%

In terms of blocked volume, gaming and telecom had to endure the highest volumes, good for over 50% of the total blocked volume in Q3. Technology, research and education, finance and healthcare mostly accounted for the remaining volume.

Research and education saw its largest volume in September, gaming accounted for most of the volume in August, and July's attack volume was mostly directed at telecom and technology.

In terms of blocked volume, gaming and telecom had to endure the highest volumes, good for over 50% of the total blocked volume in Q3. Technology, research and education, finance and healthcare mostly accounted for the remaining volume.

Research and education saw its largest volume in September, gaming accounted for most of the volume in August, and July's attack volume was mostly directed at telecom and technology.

FIGURE 16:
Top attacked industries in Q3 of 2021, normalized per customer

Blocked attacks by industry (normalized) in Q3 of 2021
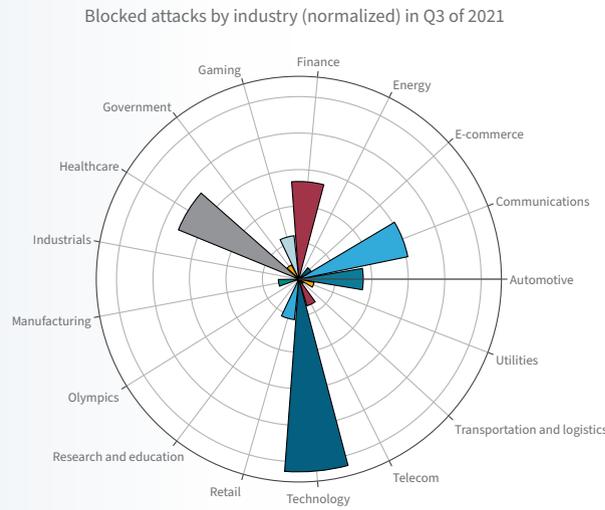
Top attacked industries (normalized) in Q3 of 2021
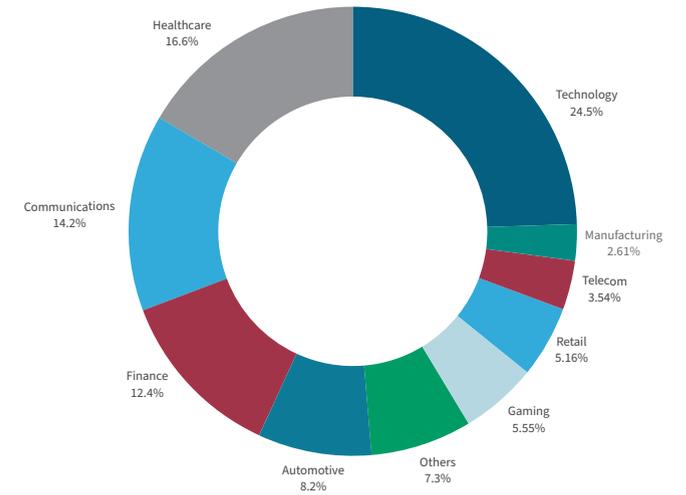
FIGURE 17:
Volume by industry for Q3 of 2021

Q3 volume by industry
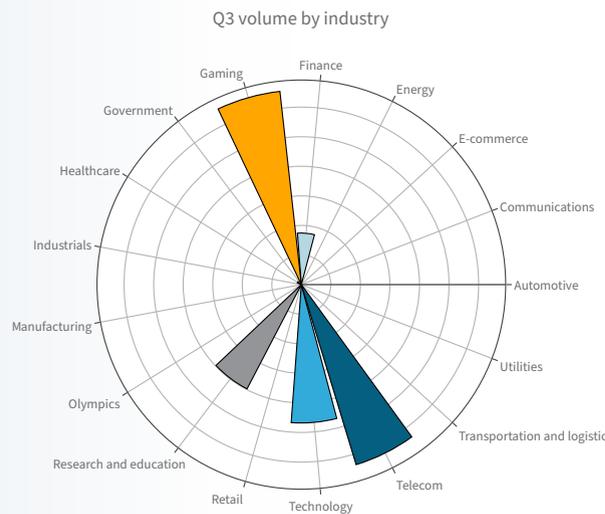
Q3 volume by industry
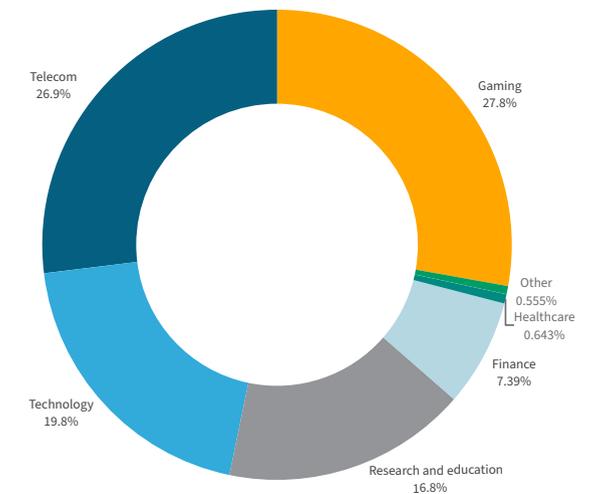
ATTACK VECTORS AND APPLICATIONS

UDP Fragmentation attacks still accounted for the bulk of the attack volume in Q3. SSL-ClearText traffic saw an increase in attack volume and became more significant than the volume originating from UDP Floods. TCP-based attacks accounted for a significant portion of the traffic in Q3, which is atypical given the amplification opportunity provided by UDP services and the typically small size involved in spoofed TCP-based attacks that are limited by the three-way-handshake requirements built into the protocol.

HTTPS, HTTP and SMTP applications accounted for most of the attack volume per customer.

FIGURE 18:
Top attack vectors by volume, normalized per customer

Top attack vectors by volume (normalized) in Q3 of 2021
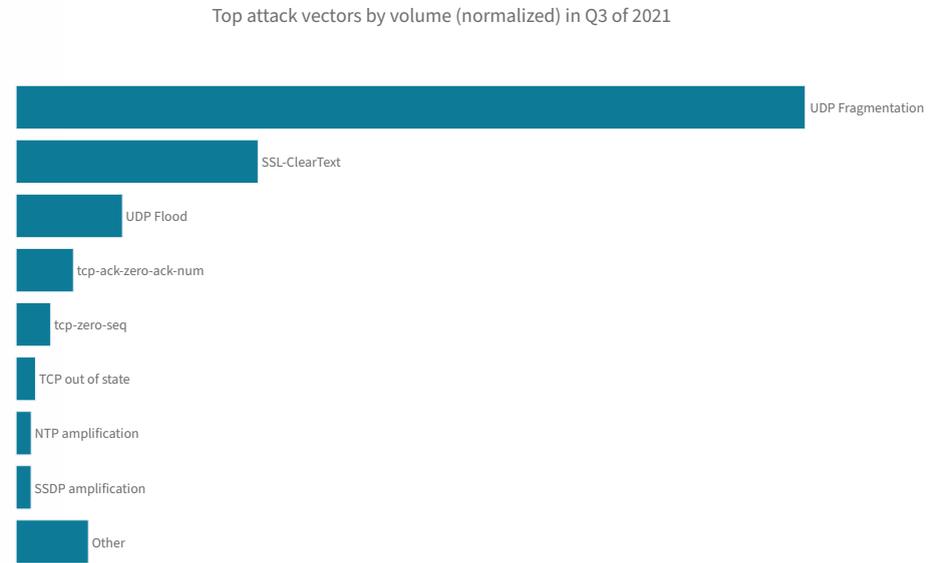
FIGURE 19:
Top applications by volume, normalized per customer

Q3 top applications by volume (normalized)
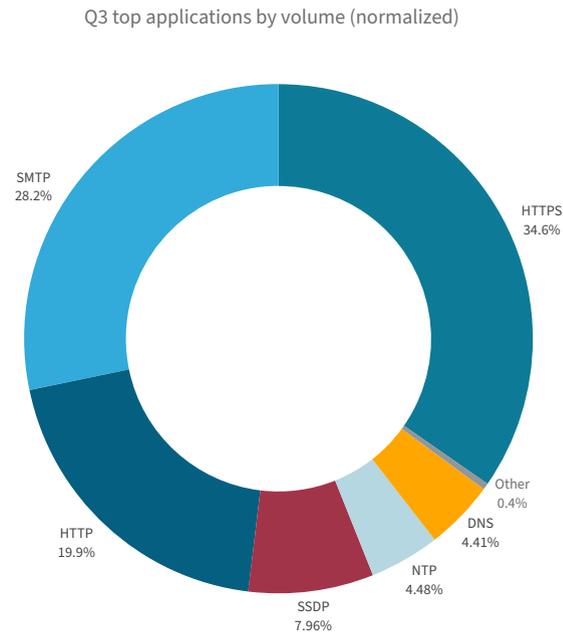
In Q1 and Q2, UDP-based attacks accounted for 99% and 97%, respectively, of the total blocked volume. In Q3, however, UDP accounted for a more moderate 88.4% of the total attack volume while TCP represented over 10% of the total blocked volume.

The low total volume in Q3 compared to the other quarters in 2021 and the shift in targeted applications, protocols and attack vectors illustrate the shift in DDoS tactics from saturation-based flooding to application-level attacks.

AMPLIFICATION ATTACK VECTORS

On average, NTP, SSDP, CLDAP and DNS were the most-used amplification attack vectors in Q3.

ARM and SSDP were the most-used amplification attacks vectors in 2020. In 2021, attackers predominantly leveraged NTP in Q1 and Q2 while DNS, NTP, SSDP and CLDAP were the preferred amplification attack vectors targeting customers in Q3.

FIGURE 20:
Protocols by volume, normalized per customer

Protocols by volume (normalized) in Q3 of 2021

GGP 0.00431%
IP 0.0923%
ICMP 1.08%
TCP 10.4%
UDP 88.4%
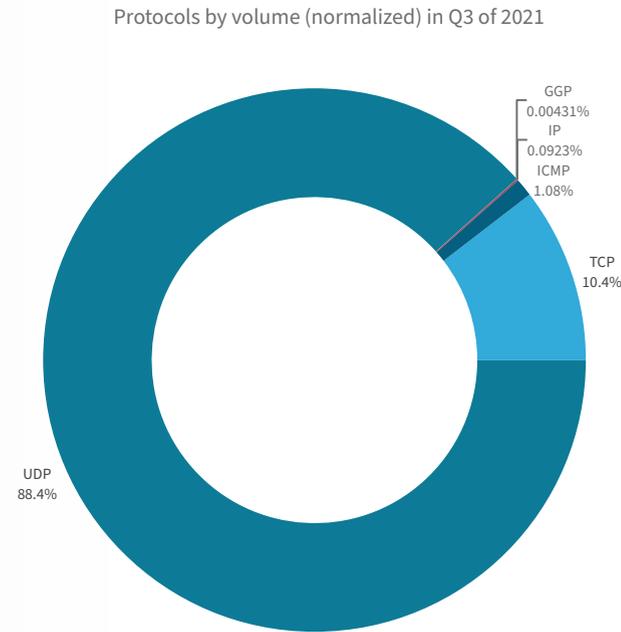
FIGURE 21:
Top amplification volumes, normalized per customer

Top amplification volumes (normalized) in Q3 of 2021

NTP 32.4%
SSDP 32.2%
CLDAP 16.6%
DNS 15%
Chargen 0.0000529%
Memcached 0.00332%
NXNS 0.0201%
COAP 0.154%
DHDiscover 0.6%
ARMS 1.1%
SNMP 2.02%
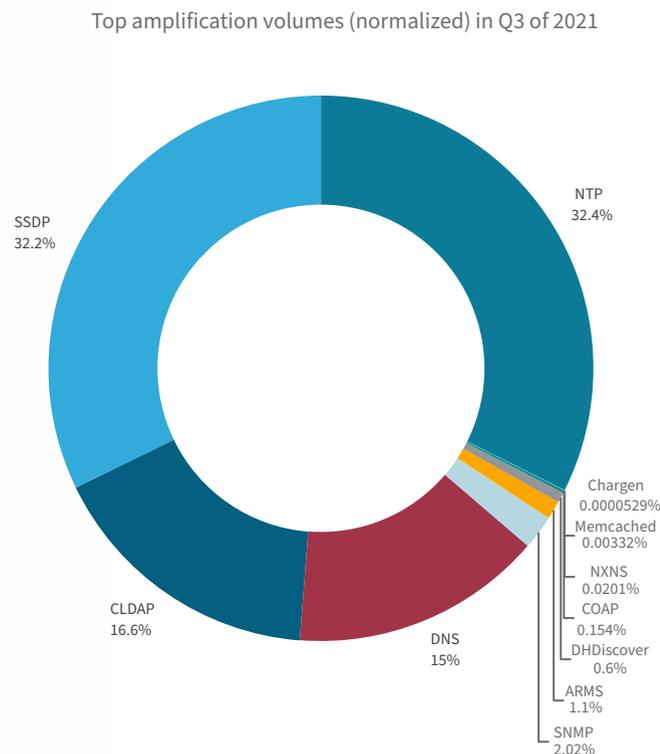
NETWORK INTRUSIONS

The number of intrusion attacks is typically larger than the number of denial-of-service (DoS) attacks. This is no different compared to other periods and, across the board, DoS accounts for 25% of blocked events while intrusions represent about 75% of those events.

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

FIGURE 22:
Top amplification attack vectors by volume over time, normalized per customer

Top amplification attack vectors by volume (normalized)
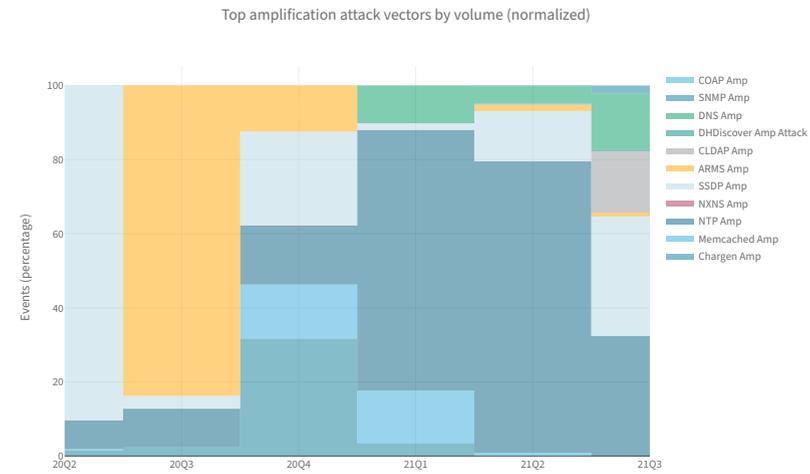
FIGURE 23:
Blocked events by attack categories

Q3 attack categories

DoS
27.8%
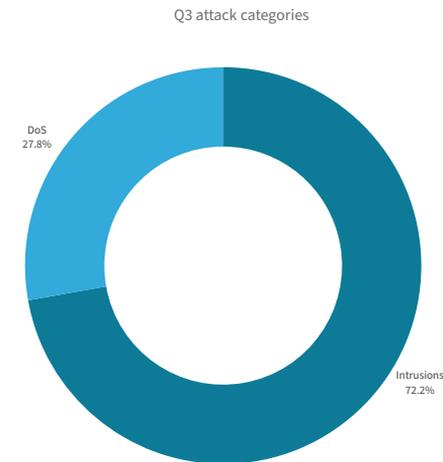
Intrusions
72.2%

FIGURE 24:
Top blocked network intrusions

Q3 top network intrusions

SIP-Scanner-SIPVicious
HTTP-Reply-MS-IE-MalfrmdBMPBO
SMTP-MS-Excel-BO
DNS-named-version-attempt
HTTP-MISC-ZMEU-SCANNER
HTTP-Reply-MS-Excel-IMData-BO
HTTP-Reply-Windows-ANI-LAI-BO
Web-etc/passwd-Dir-Traversal
HTTP-Reply-MS-Excel-SubPar-IO
HTTP-APACHE-TOMCAT-GET-OPEN

RADWARE ID	CLASSIFICATION	COMMON VULNERABILITIES AND EXPOSURES (CVE)
SIP-Scanner-SIPVicious	Scanning	–

SIPVicious – SIP information gathering and scanning tool. It detects SIP devices and identifies active extensions on a PBX and the existence of known vulnerabilities.

RADWARE ID	CLASSIFICATION	COMMON VULNERABILITIES AND EXPOSURES (CVE)
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

RADWARE ID	CLASSIFICATION	COMMON VULNERABILITIES AND EXPOSURES (CVE)
SMTP-MS-Excel-BO	Buffer overflow	CVE-2007-3890

Microsoft Excel Workspace Index Value Memory Corruption – 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 can 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.

RADWARE ID	CLASSIFICATION	COMMON VULNERABILITIES AND EXPOSURES (CVE)
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.

RADWARE ID	CLASSIFICATION	COMMON VULNERABILITIES AND EXPOSURES (CVE)
HTTP-MISC-ZMEU-SCANNER	Scanning	–

ZmEu – A vulnerability scanner that searches for web servers that are vulnerable to attacks. It also attempts to guess passwords through brute-force methods, which may lead to DoS.

RADWARE ID	CLASSIFICATION	COMMON VULNERABILITIES AND EXPOSURES (CVE)
HTTP-Reply-MS-Excel-IMData-BO	Buffer overflow	CVE-2007-0027

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.

RADWARE ID	CLASSIFICATION	COMMON VULNERABILITIES AND EXPOSURES (CVE)
HTTP-Reply-Windows-ANI-LAI-BO	Buffer overflow	CVE-2007-0038

Windows ANI "LoadAniIcon()" – 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 LoadAniIcon() 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.

RADWARE ID	CLASSIFICATION	COMMON VULNERABILITIES AND EXPOSURES (CVE)
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 [4]. 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-MS-Excel-SubPar-IO	Buffer Overflow	CVE-2011-0097

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.

RADWARE ID	CLASSIFICATION	COMMON VULNERABILITIES AND EXPOSURES (CVE)
HTTP-APACHE-TOMCAT-GET-OPEN	Information Compromise	CVE-2018-11784

Apache Tomcat HTTP open redirection – A 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.

Q3 Drivers for DDoS

"REVIL" RANSOM DOS TARGETING VOIP TELECOMMUNICATIONS

September was not only the month when DDoS began targeting education, but it was also marked by service impacting DDoS attacks on VoIP telecommunications providers in the United Kingdom and Canada. Starting September 1, UK South Coast–based VoIP operator Voip Unlimited disclosed it was hit by a sustained and large-scale DDoS attack it believed originated from the Russian ransomware group "REvil" following what they described as a "colossal ransom demand." After 75 hours of continuous attacks, on September 3, Voip Unlimited reported a pause in malicious traffic and confirmed a few days later that they did not observe any further attacks.

At the same time, also starting September 1, the London-based Voipfone reported suffering outages on voice services, inbound and outbound calls, and SMS services. It was later confirmed to customers via e-mail that Voipfone services had been "intermittently disrupted by a DDoS attack" [5].

On September 16, a Canadian provider of telephony services, VoIP.ms, announced it became aware of issues preventing customers from accessing its website and were working toward a solution. One week later, the issue was still ongoing and was attributed to persistent aggressive DDoS attacks causing disruptions in phone calls and services [6].

Public messages exchanged on Twitter between VoIP.ms and the threat actors going by the handle @REvil92457183 provided more insights. The threat actors behind the DDoS assault went by the name "REvil," but there is no evidence they represent the same REvil ransomware gang that is known to have previously attacked prominent companies, including the world's largest meat

FIGURE 25:
Now-removed Pastebin ransom demand note to the attention of VoIP.ms

```
0.52 KB                                              raw   download   report

 1.
 2.    We hope you have not suffered too much from our small demonstration...
 3.
 4.    You now have the option of paying us 1 Bitcoin or suffer further service disruption
 5.
 6.    Payment should be made during the next hour diligently or we assign a team and make an example of ▮▮▮▮
 7.
 8.    If you do not comply promptly, we will also take out your 60+ servers without breaking a sweat
 9.    Comply or fight, we do not give a fuck as we have nothing to lose...
10.
11.    1 Bitcoin: ▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮
12.
13.
14.    REvil
15.    This is OUR Dominion
16.
```

processor, JBS. As a ransomware operator, it is not conforming to the tactics, techniques and procedures (TTPs) of REvil to perform DDoS extortion attacks. It cannot be excluded, though. It would not be the first time a criminal group is diversifying its activities.

A now-removed Pastebin note put the initial ransom demand at 1 bitcoin, or a little over US$42,000 (at the time of publishing). However, only two days after the initial demand, the @ REvil92457183 twitter account increased the demand to 100 bitcoins, or over US$4.2 million, when it messaged: "Ok, enough communication... The price for us to stop is now 100 Bitcoin into the pastebin BTC address. I am sure your customers will appreciate your 0 f…s given attitude in multiple law suits. REvil".

The actors used Twitter to expose the attacks and condemn VoIP.ms for not paying, in an attempt to make its customers and partners put pressure on the service provider to pay the ransom and rid them of service disruptions. Ransomware operators use similar pressure tactics, such as leaking new victims and sensitive data obtained from victims on their dark web PR sites.

Messages from VoIP.ms on Twitter show that the attacks were initially targeting VoIP.ms' domain name services. VoIP.ms mitigated these attacks by asking partners and customers to hardcode service IP addresses in their systems.

On October 7, Voip Unlimited again reported intermittent loss of connectivity and voice services as its engineers were working to mitigate DDoS attacks against its telephony platforms. Issues were first reported on its service status webpage in the evening on October 7. In the evening on October 8, VoIP Unlimited was still reporting disruptions caused by ongoing DDoS attacks against its customers.

FIGURE 26:
VoIP.ms disclosing attacks on domain name services through Twitter
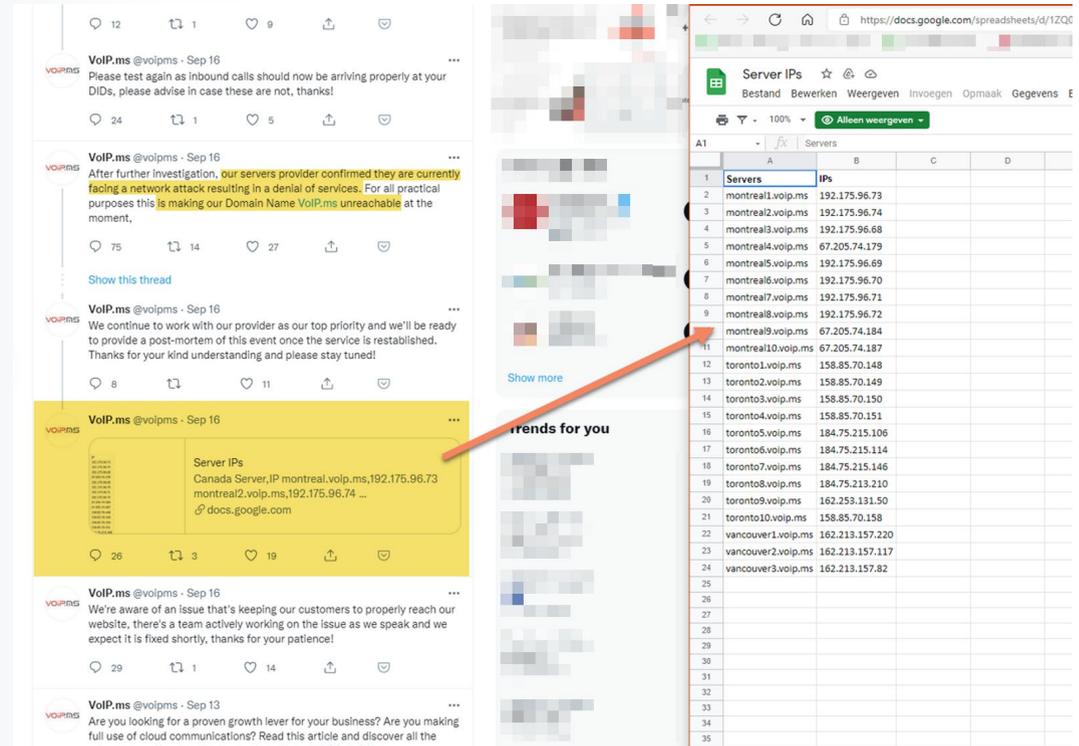
RECORD-LEVEL DDOS ATTACKS

August was a month where DDoS attack records were broken and challenged. While Radware had a rather slow month in terms of volume, other organizations were less fortunate.

On August 19, Cloudflare reported having mitigated a 17.2 million request-per-second (rps) HTTP DDoS attack [7]. The attack lasted just over 60 seconds and originated from more than 20,000 bots in 125 countries around the globe. Most bots were located in Indonesia, India, Brazil and Vietnam.

On September 9, Qrator Labs, a Russian-based DDoS attack mitigation company, published a blog about new record-level HTTP DDoS attacks targeting Yandex, a Russian multinational corporation providing internet-related products and services including transportation, search and information services, e-commerce, navigation, mobile applications and online advertising [8]. Throughout August, Yandex was faced with several high-rate rps HTTP attacks ranging from 5.2 million to 10.9 million rps; and on September 5, Yandex faced a record 21.8 million rps HTTP attack lasting about 60 seconds. According to Yandex and Qrator Labs, the attacks leveraged a HTTP/1.1 pipelining technique. They also confirmed the 56,000 originating devices of attack traffic were MikroTik devices located mostly in Brazil, Indonesia, India and Bangladesh, and they noted similarities with the geolocation of the attack sources reported in the Cloudflare attack. Yandex and Qrator Labs claim all MikroTik devices had SOCKS4 services enabled. They dubbed what they believe is a new kind of botnet, Mēris. Mēris means "Plague" in the Latvian language, and they said, "It seems appropriate and relatively close to Mirai in terms of pronunciation." "Appropriate" refers to the Latvian roots of MikroTik in this case.

MikroTik responded to the Mēris botnet and a potential new vulnerability in RouterOS, saying, "As far as we have seen, these attacks use the same routers that were compromised in 2018, when MikroTik RouterOS had a vulnerability that was quickly patched. Unfortunately, closing the vulnerability does not immediately protect these routers. If somebody got your password in 2018, just an upgrade will not help. You must also change password, recheck your firewall if it does not allow remote access to unknown parties and look for scripts that you did not create."

On October 11, Microsoft posted a blog about a 2.4Tbps DDoS attack targeting a Microsoft Azure customer in Europe in August 2021. This attack was higher than any network volumetric event previously detected on Azure. The attack's peak traffic volume exceeded the 2.3 Tbps assault on Amazon Web Services last year, though it was a smaller attack compared to the 2.54 Tbps attack Google mitigated in 2017. The attack vector was a UDP Reflection attack spanning more than 10 minutes with short-lived bursts, each ramping up in seconds to terabit volumes. In total, they reported three peaks – the first at 2.4 Tbps, the second at 0.55 Tbps, and the third at 1.7 Tbps. Each peak lasted about 60 seconds.

Web Application Attack Activity

The number of web security events blocked by the Radware Cloud WAF Service has doubled every quarter for the first three quarters of 2021. Q3 accounts for 2.1 million blocked security events per customer per quarter, or an average of 700,000 blocked security events per month per customer.

The number of blocked web application security events include only automatically detected and known vulnerabilities and exclude all custom rules potentially added to a web application policy by managed services and/or customers.

Most blocked web security events originated from the United States and Russia. India, the United Kingdom and Germany completed the top five in Q3 of 2021.

Application security events (normalized per customer) in 2021

FIGURE 27:
Quarterly blocked application security events

Q3 top offending countries
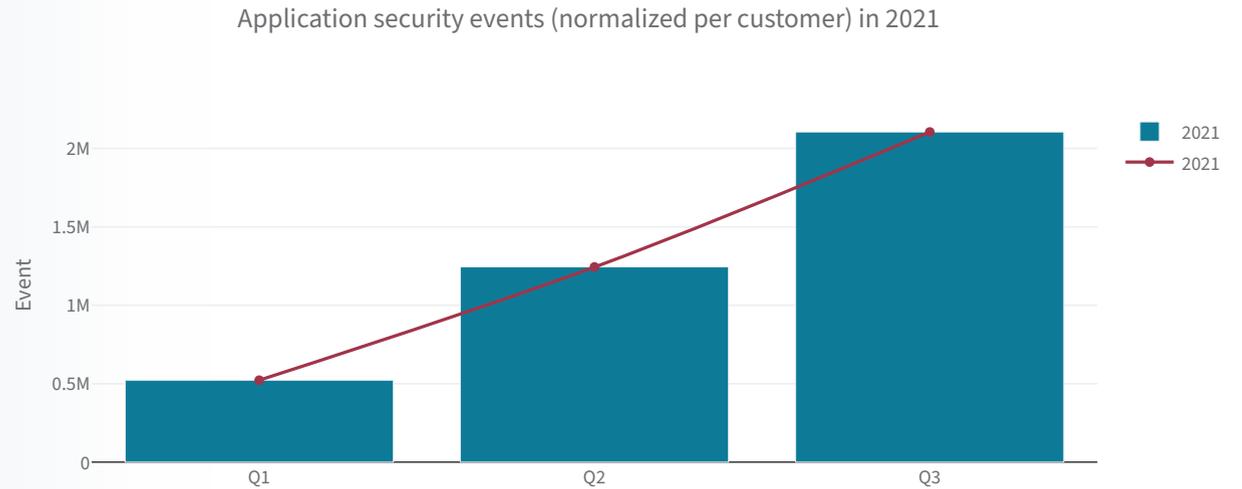
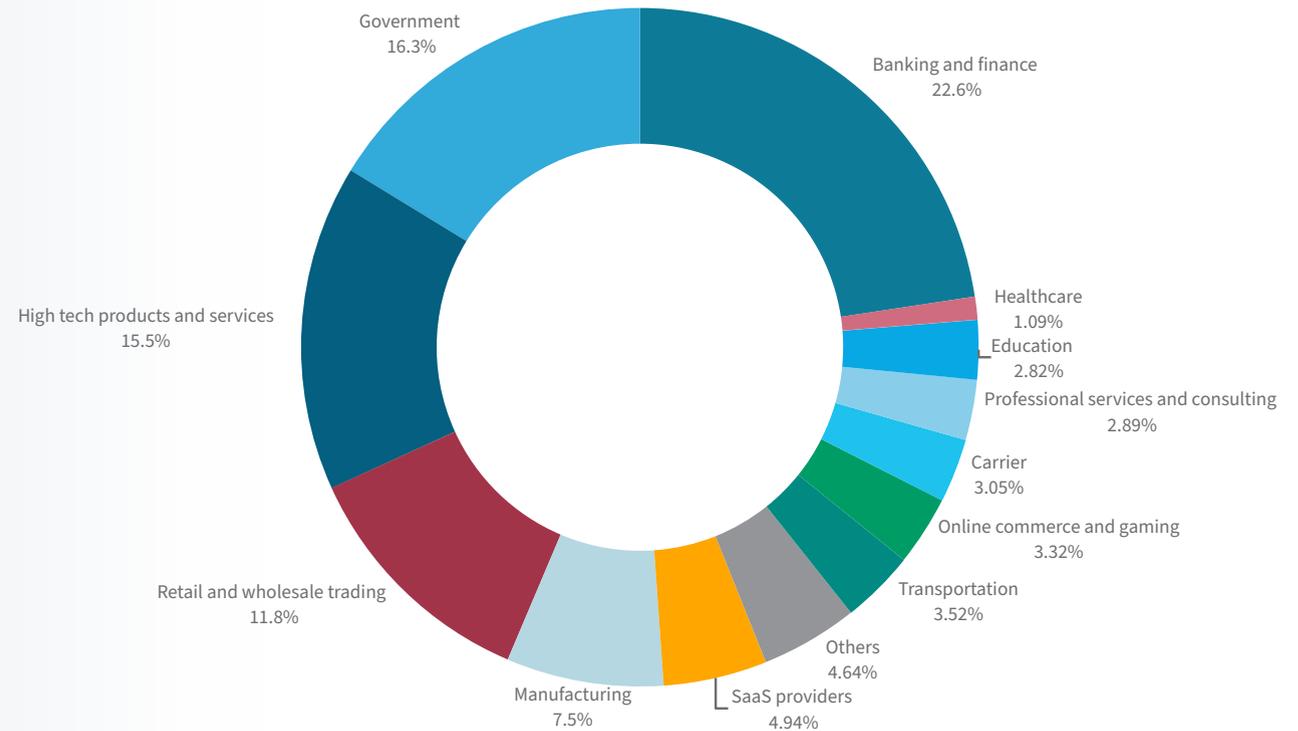**FIGURE 28:**
Top offending countries
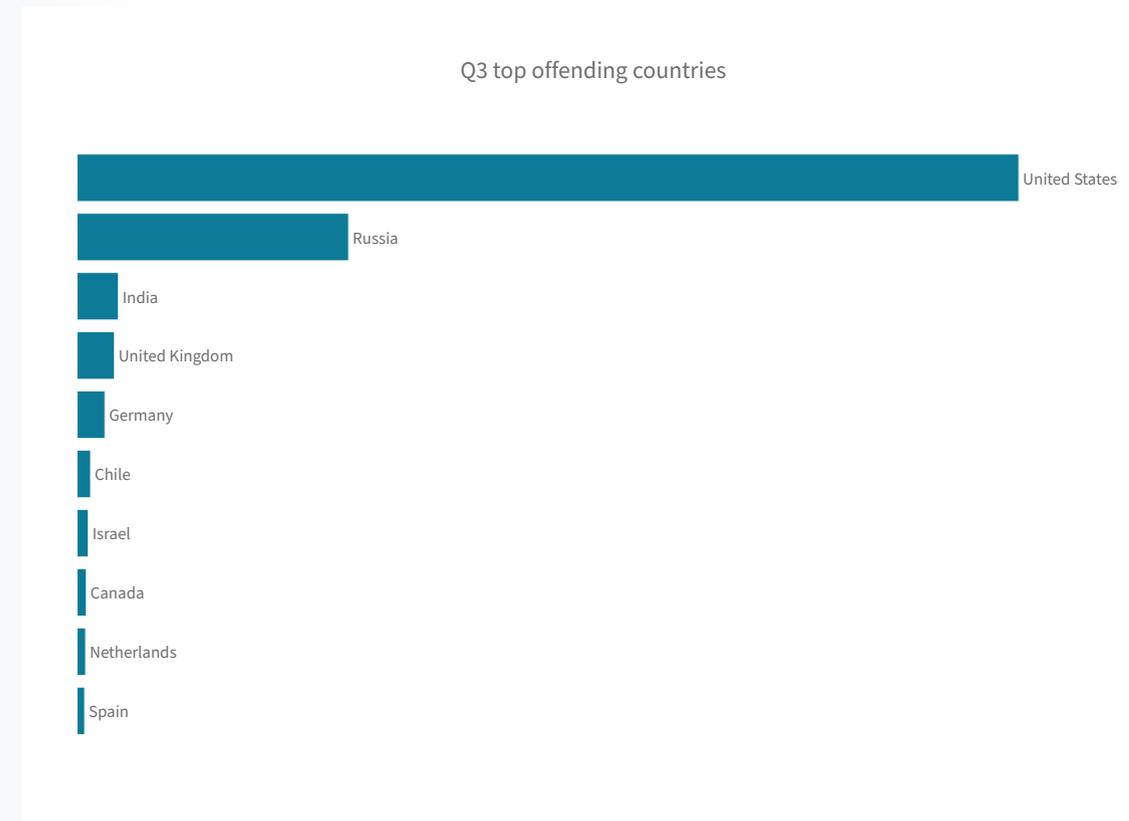
The most attacked industry was banking and finance, with 22.6% of blocked web security events, followed by government (16.3%), technology (15.5%) and retail and wholesale trading (11.8%).

The most important security violation – predictable resource location attacks – was witnessed twice as often as the second-most violation – SQL injection. Predictable resource location attacks target hidden content and functionality of web applications. By guessing common names for file directories, an attack may be able to access resources unintended for exposure. Examples of resources that might be uncovered through brute-force techniques include old backup and configuration files, web application resources yet to be published and so on. Predictable resource location attempts are covered by the OWASP 2017 Top 10[2] web application security risk "Broken Access Control," which was ranked 5th in 2017 and moved to first place in the 2021 OWASP Top 10.

FIGURE 29:
Attacks per industry, normalized per customer

Q3 attacked industries (normalized per customer)

- Government 16.3%
- Banking and finance 22.6%
- Healthcare 1.09%
- Education 2.82%
- Professional services and consulting 2.89%
- Carrier 3.05%
- Online commerce and gaming 3.32%
- Transportation 3.52%
- Others 4.64%
- SaaS providers 4.94%
- Manufacturing 7.5%
- Retail and wholesale trading 11.8%
- High tech products and services 15.5%

2.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.
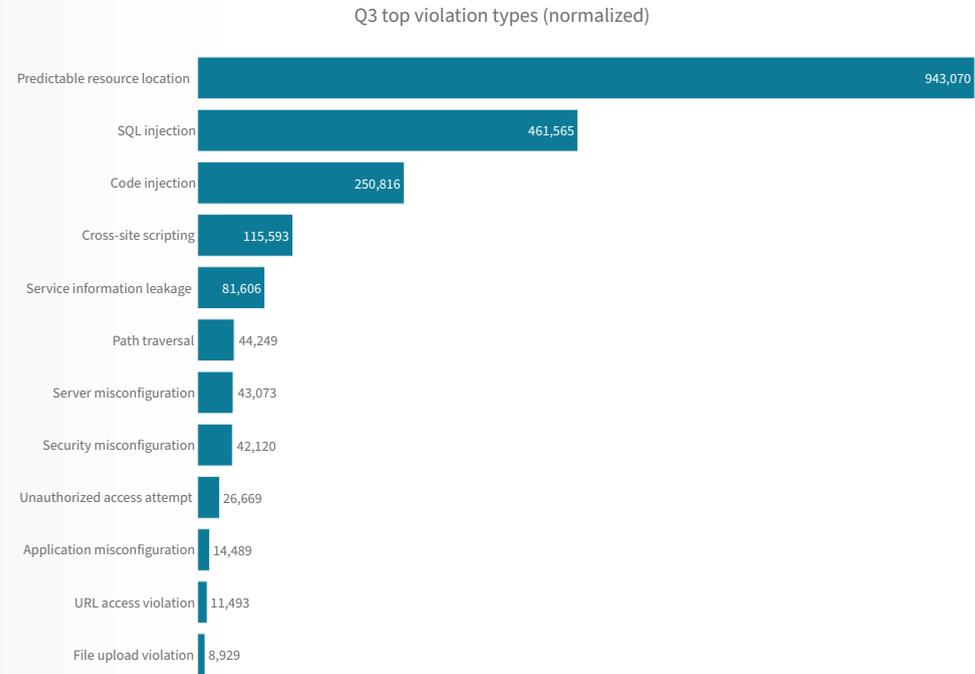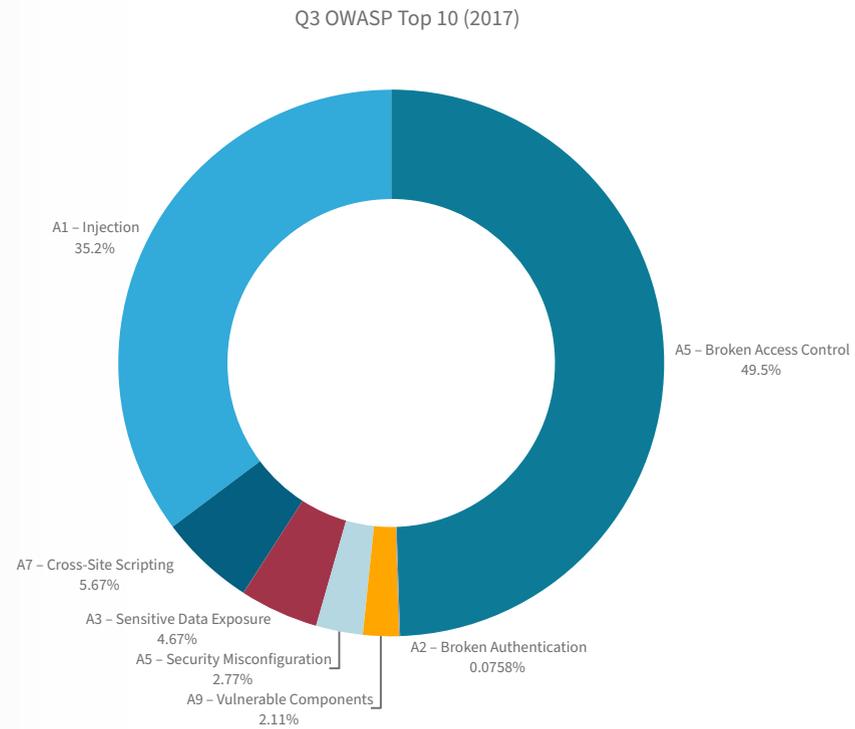
The number-one web application security risk, according to the 2017 OWASP Top 10, is injection attacks, as illustrated by the SQL injection and code injection top violation types in Q3. Cross-site script took fourth place in Q3 and corresponds to the cross-site scripting (XSS) (A7) OWASP application security risk.

FIGURE 30:
Top security violation types, normalized per customer

Q3 top violation types (normalized)

Violation type	Value
Predictable resource location	943,070
SQL injection	461,565
Code injection	250,816
Cross-site scripting	115,593
Service information leakage	81,606
Path traversal	44,249
Server misconfiguration	43,073
Security misconfiguration	42,120
Unauthorized access attempt	26,669
Application misconfiguration	14,489
URL access violation	11,493
File upload violation	8,929

FIGURE 31:
Blocked security violations by OWASP 2017 application security risks

Q3 OWASP Top 10 (2017)

A1 – Injection 35.2%
A5 – Broken Access Control 49.5%
A7 – Cross-Site Scripting 5.67%
A3 – Sensitive Data Exposure 4.67%
A5 – Security Misconfiguration 2.77%
A9 – Vulnerable Components 2.11%
A2 – Broken Authentication 0.0758%

Unsolicited Network Scanning and Attack 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, spoofed and nonspoofed scans, and spoofed[3] and nonspoofed attacks. The difference between deception network events discussed in this section and the web application and DDoS attack events in previous sections is the unsolicited nature of the event. Web application and DDoS attack events were collected from services that protect published services of organizations, backed by real applications and networks. The attackers were targeting a particular organization or a known service. 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 Radware's Global Deception Network sensors.

Q3 activity peaked at 27 million events per day, with almost 300 million events in August.

FIGURE 32:
Number of events per day recorded by Radware's Global Deception Network

Number of events

Per day
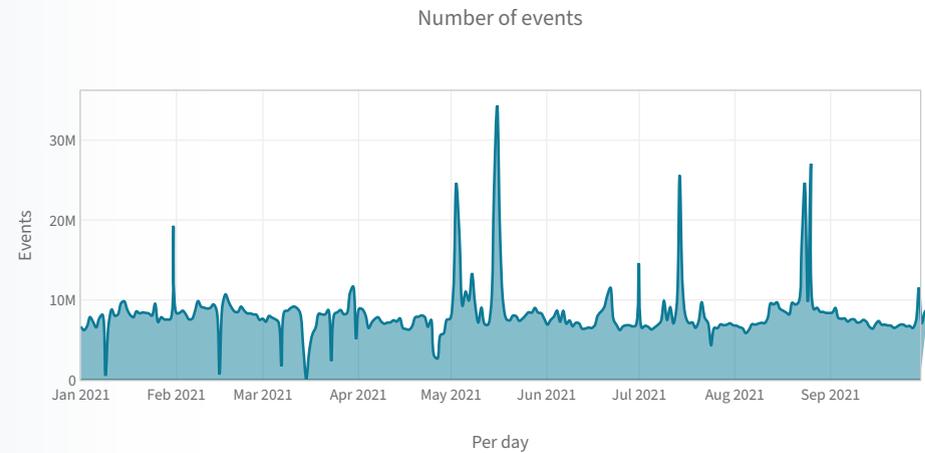
FIGURE 33:
Number of events per month recorded by Radware's Global Deception Network

Number of events

Per month

3. IP address spoofing, or IP spoofing, is the crafting of Internet Protocol (IP) packets with false source IP addresses for the purpose of impersonating another originating computing system and geolocation. (Source: Wikipedia)
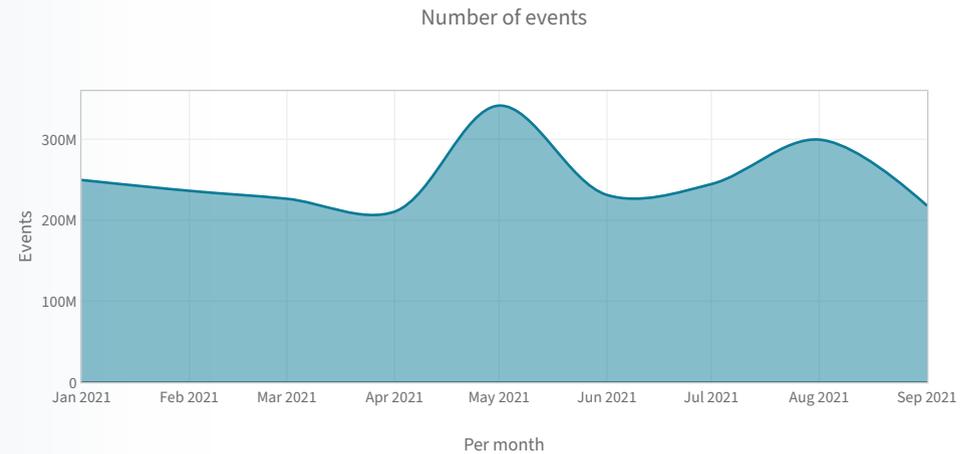
The total number of unique client or source IP addresses was at its high for the year during the summer months. However, source IP addresses can be spoofed; and to authoritatively determine an attacking IP or device, only IP addresses that were involved in a full TCP three-way handshake should be trusted. The red areas in unique IP charts indicate the number of unique IPs verified by a three-way handshake. The blue areas comprise the total number of unique IP addresses, spoofed and nonspoofed.

FIGURE 34:
Number of unique IPs per day registered by Radware's Global Deception Network
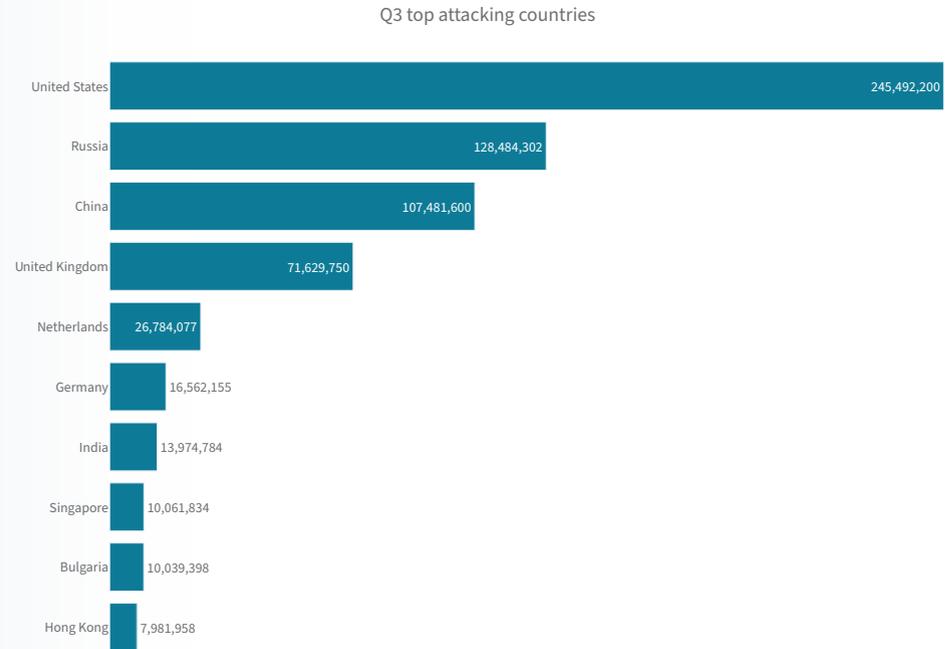
Number of unique IPs

Per day
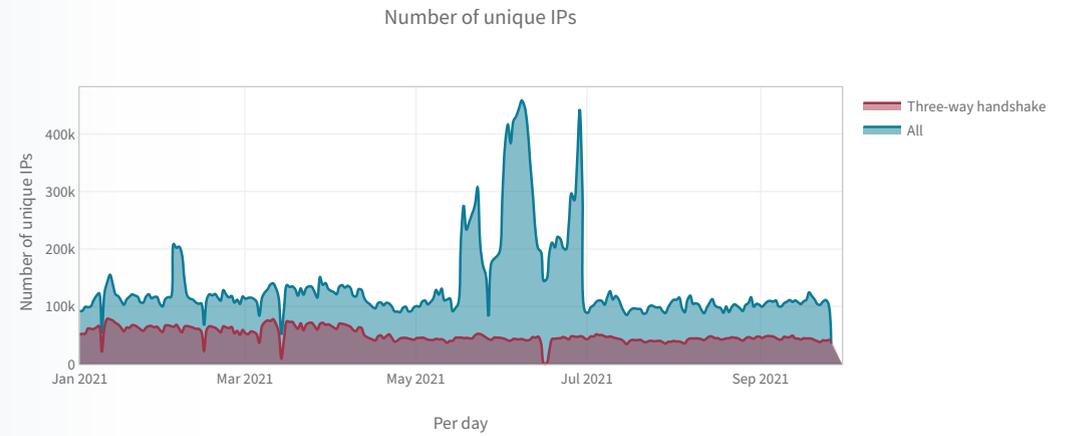
FIGURE 35:
Number of unique IPs per month registered by Radware's Global Deception Network
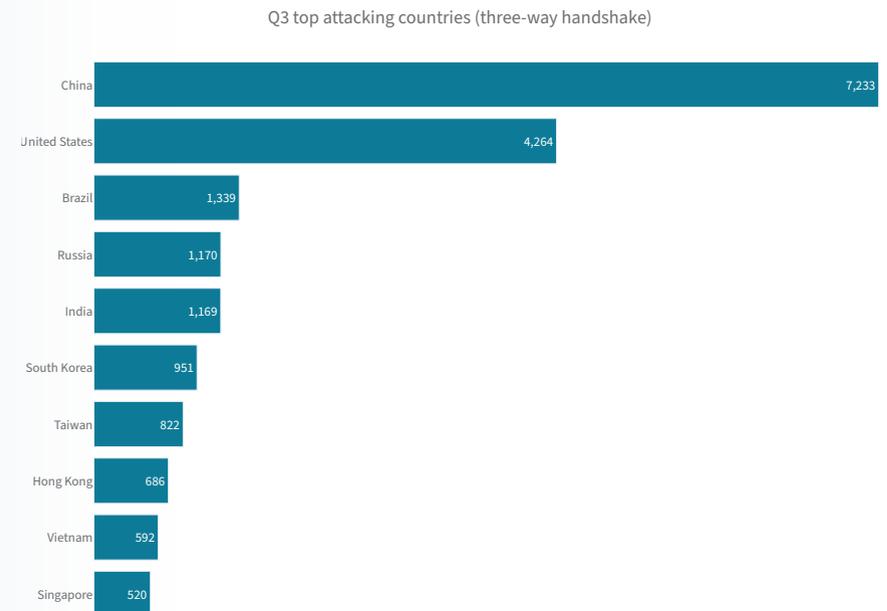
Number of unique IPs

Per month
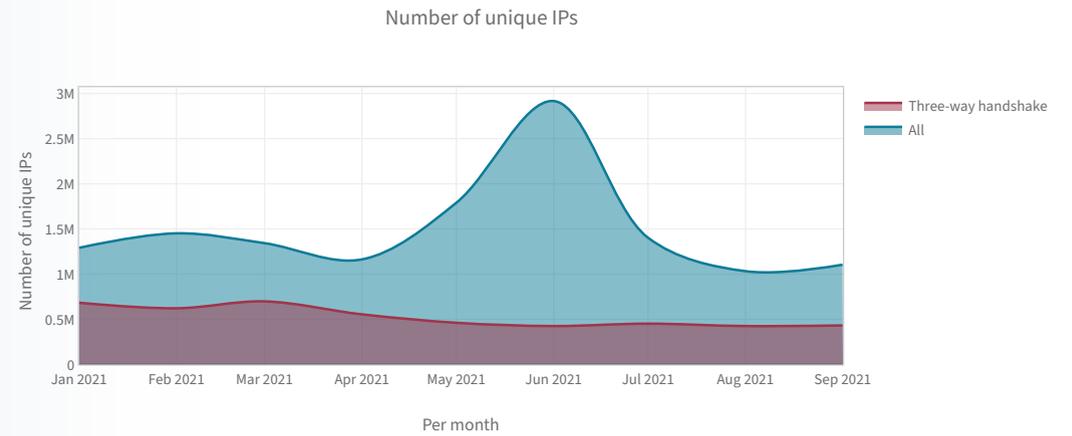
ATTACKING COUNTRIES

The top attacking countries in Q3 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. When considering attack events with verified TCP three-way handshakes only, the top attacking countries look quite different: China, the United States, Brazil, Russia and India.

Q3 top attacking countries

FIGURE 36:
Top attacking countries recorded by Radware's Global Deception Network

Country	Value
United States	245,492,200
Russia	128,484,302
China	107,481,600
United Kingdom	71,629,750
Netherlands	26,784,077
Germany	16,562,155
India	13,974,784
Singapore	10,061,834
Bulgaria	10,039,398
Hong Kong	7,981,958

Q3 top attacking countries (three-way handshake)

FIGURE 37:
Top attacking countries based on verified three-way handshake IP addresses

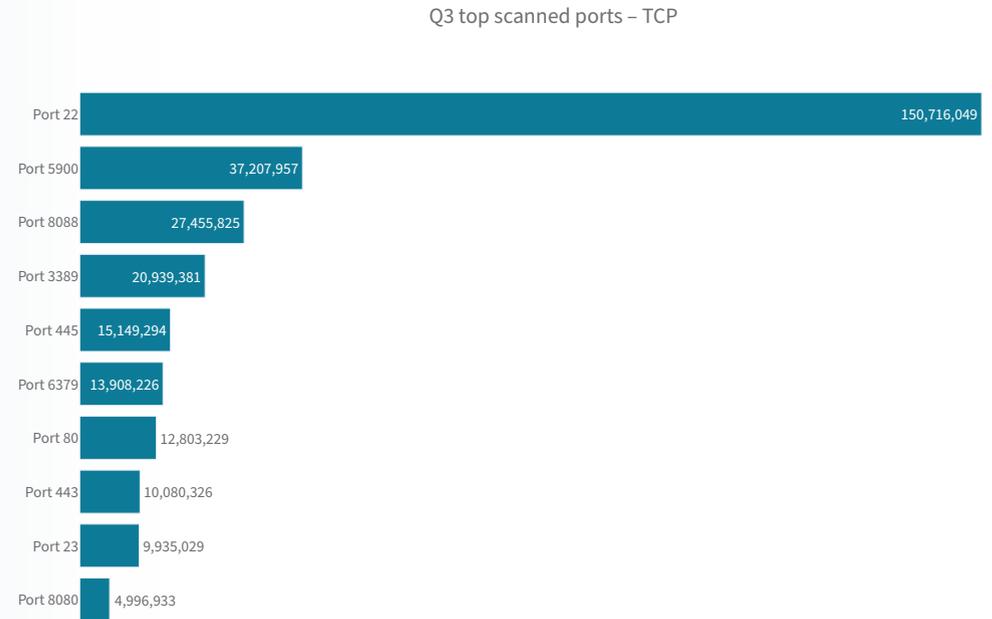| Country	Value
China	7,233
United States	4,264
Brazil	1,339
Russia	1,170
India	1,169
South Korea	951
Taiwan	822
Hong Kong	686
Vietnam	592
Singapore	520

SCANNED AND ATTACKED PORTS

For TCP services, the most scanned and attacked service was SSH on port 22, followed by VNC[4] on port 5900, HTTP on port 8088, RDP[5] on port 3389 and HTTPS on port 445.

Port 6379 is used by Redis, an open source (BSD licensed), in-memory data structure store used as a database, cache and message broker. In July, a remote command execution (RCE) 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 [1].

FIGURE 38:
Top scanned and attacked TCP ports

Q3 top scanned ports – TCP

Port	Count
Port 22	150,716,049
Port 5900	37,207,957
Port 8088	27,455,825
Port 3389	20,939,381
Port 445	15,149,294
Port 6379	13,908,226
Port 80	12,803,229
Port 443	10,080,326
Port 23	9,935,029
Port 8080	4,996,933

4. 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)

5. 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)

Telnet on port 23; HTTP on port 8088, 8080 and 80; HTTPS on port 445 and 443 remain among the top exploited TCP ports for Q3. These are typically abused by IoT botnets, including many of the Mirai variants, that continue to wreak havoc on the internet through DDoS attacks and put IoT devices such as IP cameras and routers and modems at risk. While Telnet was a Mirai favorite for a long time, the events on SSH surpassed Telnet by more than 15 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 crypto mining, leverage the foothold as jump host to anonymize targeted attacks or leverage device connectivity to perform DDoS attacks.

UDP-based services were targeted primarily through port 5060, which is used by many SIP-based VoIP phones and providers. Considering the ransom DoS attack activity targeting VoIP providers in the second half of the quarter, this activity might reflect the correlation of discovery activity by actors. Vulnerabilities in VoIP services can also be abused for initial access and move laterally inside organizations' networks.

UDP port 123 (NTP), 11211 (Memcached), 389 (LDAP), 161 (SNMP) and 1900 (SSDP, UPnP) are among the most leveraged protocols for DDoS amplification attacks. Many black and white hats are continuously scanning and cataloging the complete Internet IPv4 range to abuse (black hat) or assess the risk in the threat landscape (white hat).

FIGURE 39:
Most-attempted usernames in account takeover and Brute Force attacks on SSH
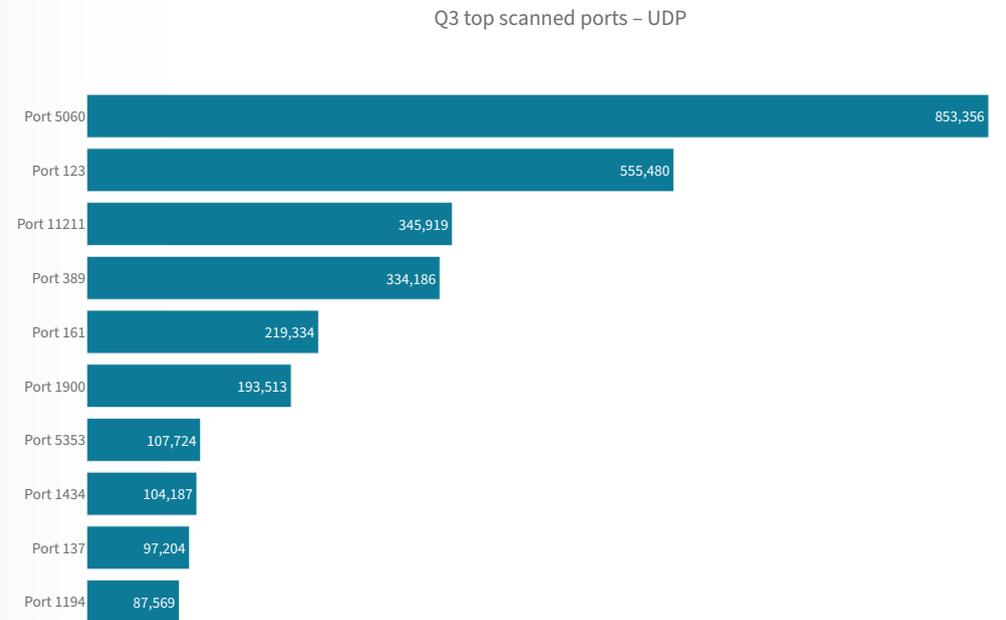
Q3 top SSH usernames
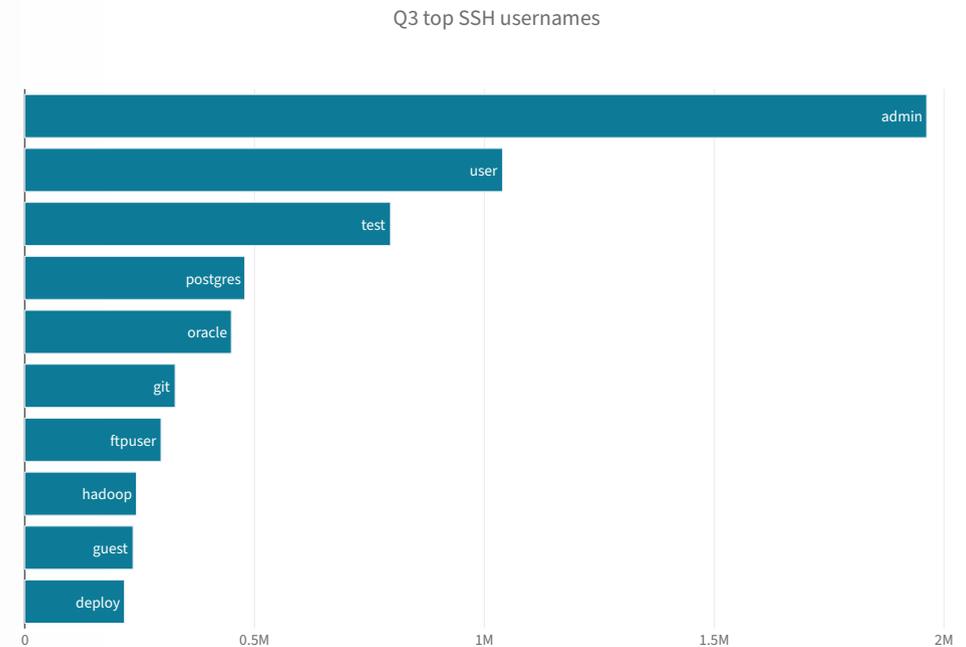
Port 5353 is typically used as an alternative for DNS (53/UDP) by small organizations or consumer networks that can expose only higher port ranges. The low port ranges (1–1024) consist of registered port numbers and are typically reserved and blocked by many ISPs on consumer internet connections. DNS services running on port 5353 are good targets for discovering misconfigured DNS services that allow amplification and reflection volumetric DDoS attacks.

Port 1434/UDP is used by the Microsoft SQL Server database management system monitor and known for a remote code execution vulnerability (CVE-2003-0353) and the W32.Spybot. Worm that spread through Microsoft SQL Server 2000 and Microsoft Desktop Engine 2000 using port 1434/UDP.

Port 137/UDP is used by the NETBIOS Name Service.
Port 1194/UDP is used by OpenVPN servers.

FIGURE 40:
Top scanned and attacked UDP ports

Q3 top scanned ports – UDP

Port	Value
Port 5060	853,356
Port 123	555,480
Port 11211	345,919
Port 389	334,186
Port 161	219,334
Port 1900	193,513
Port 5353	107,724
Port 1434	104,187
Port 137	97,204
Port 1194	87,569

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 collection 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 behind the exposed ports. The top URIs need to be interpreted as the top services and applications that are targeted by actors that are randomly scanning and exploiting these services and applications. Typically, the URI will conform to known and disclosed vulnerabilities. "/ws/v1/cluster/app/new-application" is part of a known vulnerability used to exploit Hadoop YARN services and schedule arbitrary workloads on Hadoop clusters [2]. This is an exploit seen leveraged by many cryptojacking campaigns that try to use capable on-premise or cloud-hosted Hadoop clusters of enterprises and research institutions illegitimately [3].

FIGURE 41:
Top unsolicited URI attempts for web services

Q3 top URIs

URI	Count
/	7,508,524
/ws/v1/cluster/apps/new-application	4,174,019
login.cgi	435,531
/v1.16/version	360,388
/.env	269,538
/ftptest.cgi?loginuse=&loginpas=	215,336
*	172,694
icanhazip.com:443	144,220
/manager/html	133,612
/boaform/admin/formLogin	94,516
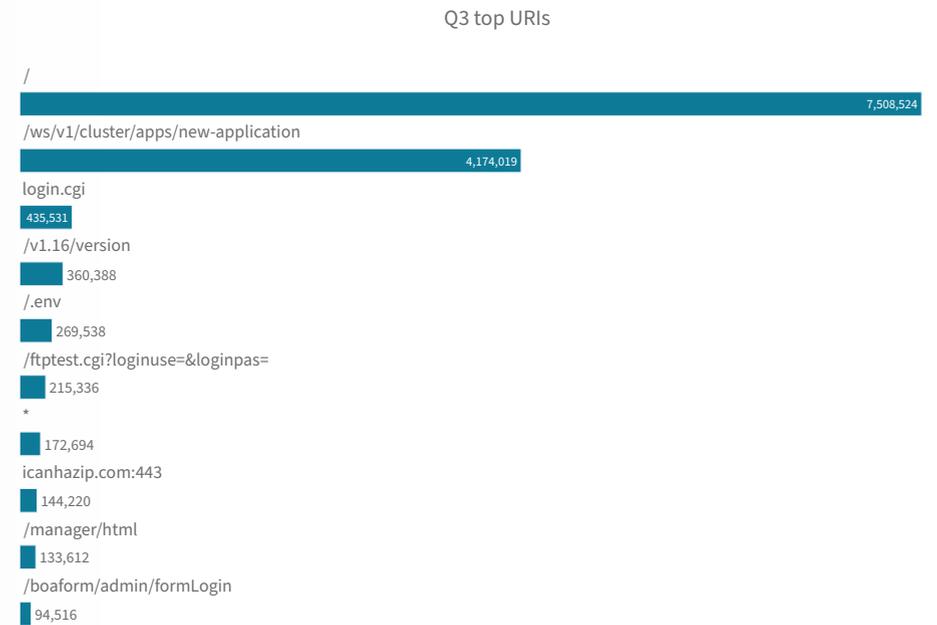
In HTTP, the User-Agent string is often used for content negotiation, where the origin server selects suitable content or operating parameters for the response. For example, the User-Agent string might be used by a web server to choose variants based on the known capabilities of a particular version of client software. 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." [9]

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 changing the user agent to known legitimate values. The original Mirai, for example, used five different user-agent headers to impersonate a browser client while performing HTTP GET attacks.

Newer Mirai variants randomly leverage several-hundred known legitimate user-agent headers to impersonate real browser agents.

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.

```c
/* User agent strings */

#define TABLE_HTTP_ONE    47  /* "Mozilla/5.0 (Windows NT 10.0; WOW64) AppleWebKit/537.36
(KHTML, like Gecko) Chrome/51.0.2704.103 Safari/537.36" */

#define TABLE_HTTP_TWO   48  /* "Mozilla/5.0 (Windows NT 10.0; WOW64) AppleWebKit/537.36
(KHTML, like Gecko) Chrome/52.0.2743.116 Safari/537.36" */

#define TABLE_HTTP_THREE 49  /* "Mozilla/5.0 (Windows NT 6.1; WOW64) AppleWebKit/537.36
(KHTML, like Gecko) Chrome/51.0.2704.103 Safari/537.36" */

#define TABLE_HTTP_FOUR  50  /* "Mozilla/5.0 (Windows NT 6.1; WOW64) AppleWebKit/537.36
(KHTML, like Gecko) Chrome/52.0.2743.116 Safari/537.36" */

#define TABLE_HTTP_FIVE  51  /* "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11_6)
AppleWebKit/601.7.7 (KHTML, like Gecko) Version/9.1.2 Safari/601.7.7" */
```
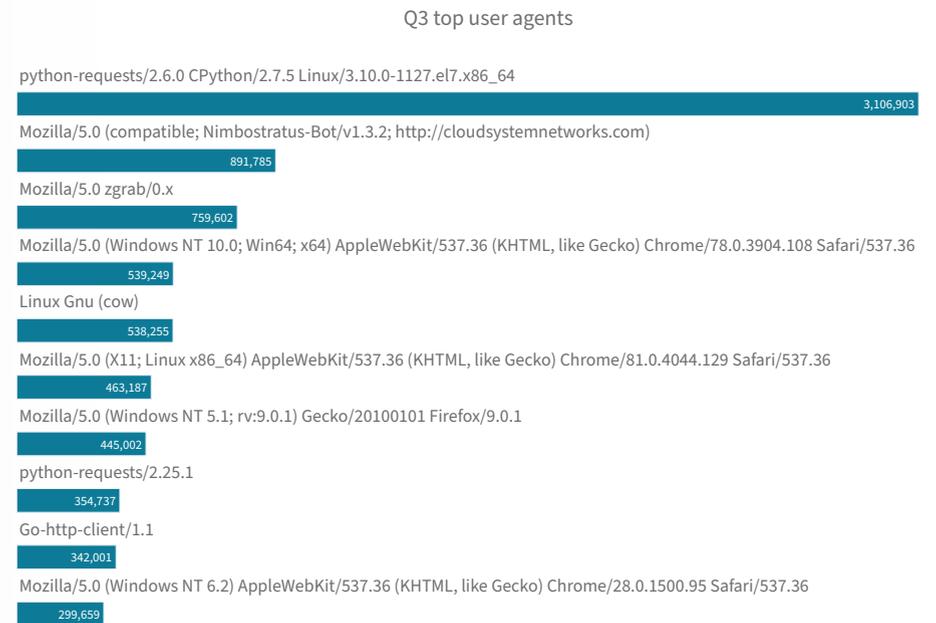
Some web crawlers and robots use the user agent to identify themselves. Websites can use a "robots.txt" file to regulate which search engine crawlers have access to which parts of the website. The "robots.txt" is a noncompulsory solution that relies completely on the crawler or robot. Needless to say, malicious bots will ignore the "robots.txt" entries and will crawl and scrape at their leisure. The "Nimbostratus-bot" is considered a legitimate bot, and Cloud System Networks leverages the user agent to make its intentions clear by adding a URL to its homepage that explains the rationale behind its activity.

Not all web service vulnerabilities can be exploited without authenticating. Some web services have widely used default or some have even hardcoded secret credentials to protect access from unauthorized users or devices. The top username and password credential pair leveraged in Q3 was "admin:admin", followed closely by "admin:123456" and "root:1234". These are universally agreed to be the worst and most abused credentials that provide a good amount of access to unauthorized devices.

"root:icatch99" and "report/8Jg0SR8K50" are two hardcoded credentials in digital video recorders (DVRs) from vendor LILIN that were publicly disclosed in March 2020 [10]. DVRs are ubiquitous in the IoT landscape, as are the security cameras that feed them.

FIGURE 42:
Top user-agent header values used for unsolicited web service requests

Q3 top user agents

python-requests/2.6.0 CPython/2.7.5 Linux/3.10.0-1127.el7.x86_64
3,106,903

Mozilla/5.0 (compatible; Nimbostratus-Bot/v1.3.2; http://cloudsystemnetworks.com)
891,785

Mozilla/5.0 zgrab/0.x
759,602

Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/78.0.3904.108 Safari/537.36
539,249

Linux Gnu (cow)
538,255

Mozilla/5.0 (X11; Linux x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/81.0.4044.129 Safari/537.36
463,187

Mozilla/5.0 (Windows NT 5.1; rv:9.0.1) Gecko/20100101 Firefox/9.0.1
445,002

python-requests/2.25.1
354,737

Go-http-client/1.1
342,001

Mozilla/5.0 (Windows NT 6.2) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/28.0.1500.95 Safari/537.36
299,659

FIGURE 43:
HTTP credentials used by web service attacks as recorded by Radware's Global Deception Network

Q3 top HTTP credentials

admin:admin
admin:123456
root:1234
root:icatch99
report:8Jg0SR8K50
admin:pass
admin:
admin:password
admin:1234567890
admin:123

0   5k   10k   15k   20k   25k   30k   35k   40k
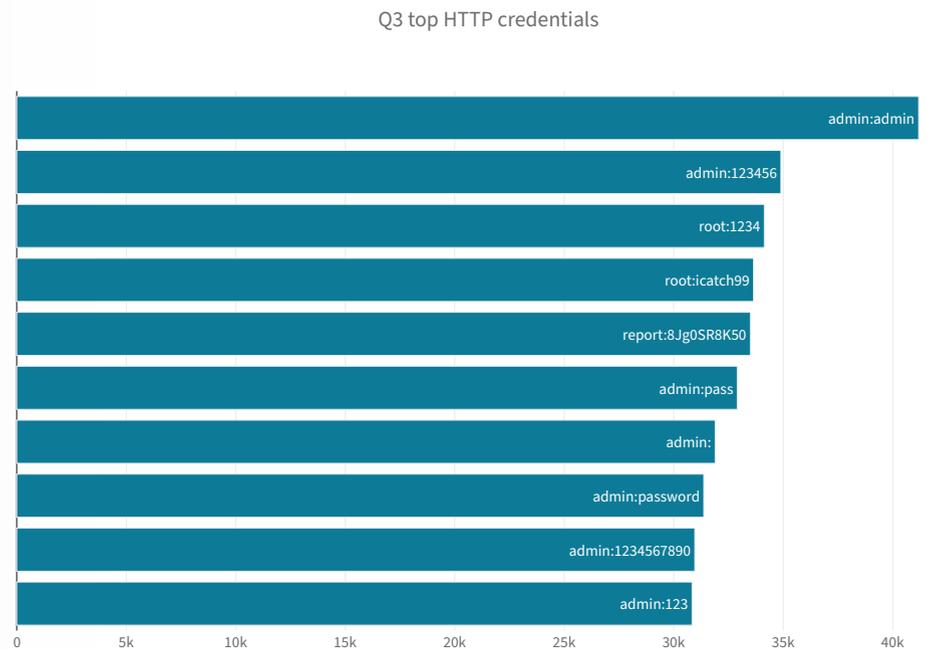
Conclusion and Key Takeaways

It is hard to draw conclusions based on short periods in an industry that is in a constant state of flux and is still adapting security measures to address changes due to the pandemic. Attacks did not regress, and overall volumes were not great. This can be partly explained by the shift to more insidious and application-level attacks.

For their ransom campaigns, actors favored targets that were not immune to their assaults and where they were able to impact services. Are they becoming more selective? Radware raised this point earlier in the year, when it reported ransom DoS attacks on targets that were not protected by always-on cloud solutions.

The record large-scale DDoS attacks were hit-and-run assaults. These might have been tests of capability or probing the protections of certain providers, or even a demonstration of capabilities and a precursor of what is yet to come. Chatter on underground forums and theories being discussed by the media do not provide a clear understanding of the objectives and tools leveraged by the actors behind those colossal assaults.

Regarding web application attacks, the number of blocked web application security events has doubled almost every quarter this year. Q3 accounts for 2.1 million blocked application security events per customer, or an average of 700,000 blocked security events per month per customer.

Almost half of the web application attacks were predictable resource location attacks. The second top security violation blocked by our web application security services was SQL injection attacks, followed by code injection attacks and cross-site scripting attacks. The top violations reported in Q3 are aligned with the top web application security risks published by the OWASP Foundation in their 2017 and 2021 OWASP Top 10 lists.

Finally, network scanning and attack activity was marked by opportunistic and random scanning that constitutes a large part of the vulnerability and exploit threat landscape. Malicious actors continuously leverage old and freshly disclosed vulnerabilities such as remote command execution and command injection exploits that are easy to integrate into existing malware and exploit tools. The objectives behind the attack activity are governed by cryptojacking, discovery of amplification and reflection services for volumetric DDoS attacks, acquiring a foothold to perform lateral movement and privilege escalation and ultimately drop backdoors or ransomware. They are also able to abuse services and devices as jump hosts or anonymous proxy and port forwarders for targeted attacks.

References

[1] 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.

[2] 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.

[3] Radware, "Demonbot," October 25, 2018. [Online]. Available: www.radware.com/security/ddos-threats-attacks/threat-advisories-attack-reports/demonbot.

[4] N. Biasini, "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.

[5] 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.

[6] A. Sharma, "Phone Calls Disrupted by Ongoing DDoS Cyberattack on VOIP.ms," Ars Technica, September 22, 2021. [Online]. Available: https://arstechnica.com/gadgets/2021/09/canadian-voip-provider-hit-by-ddos-attack-phone-calls-disrupted.

[7] 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.

[8] Qrator, "Mēris Botnet, Climbing to the Record," September 9, 2021. [Online]. Available: https://blog.qrator.net/en/meris-botnet-climbing-to-the-record_142.

[9] Wikipedia, "User Agent," September 30, 2021. [Online]. Available: https://en.wikipedia.org/wiki/User_agent.

[10] 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.

List of Figures

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.