[Volume 2 / Number 3]

akamai's [state of the internet] / security

Q3 2015 report

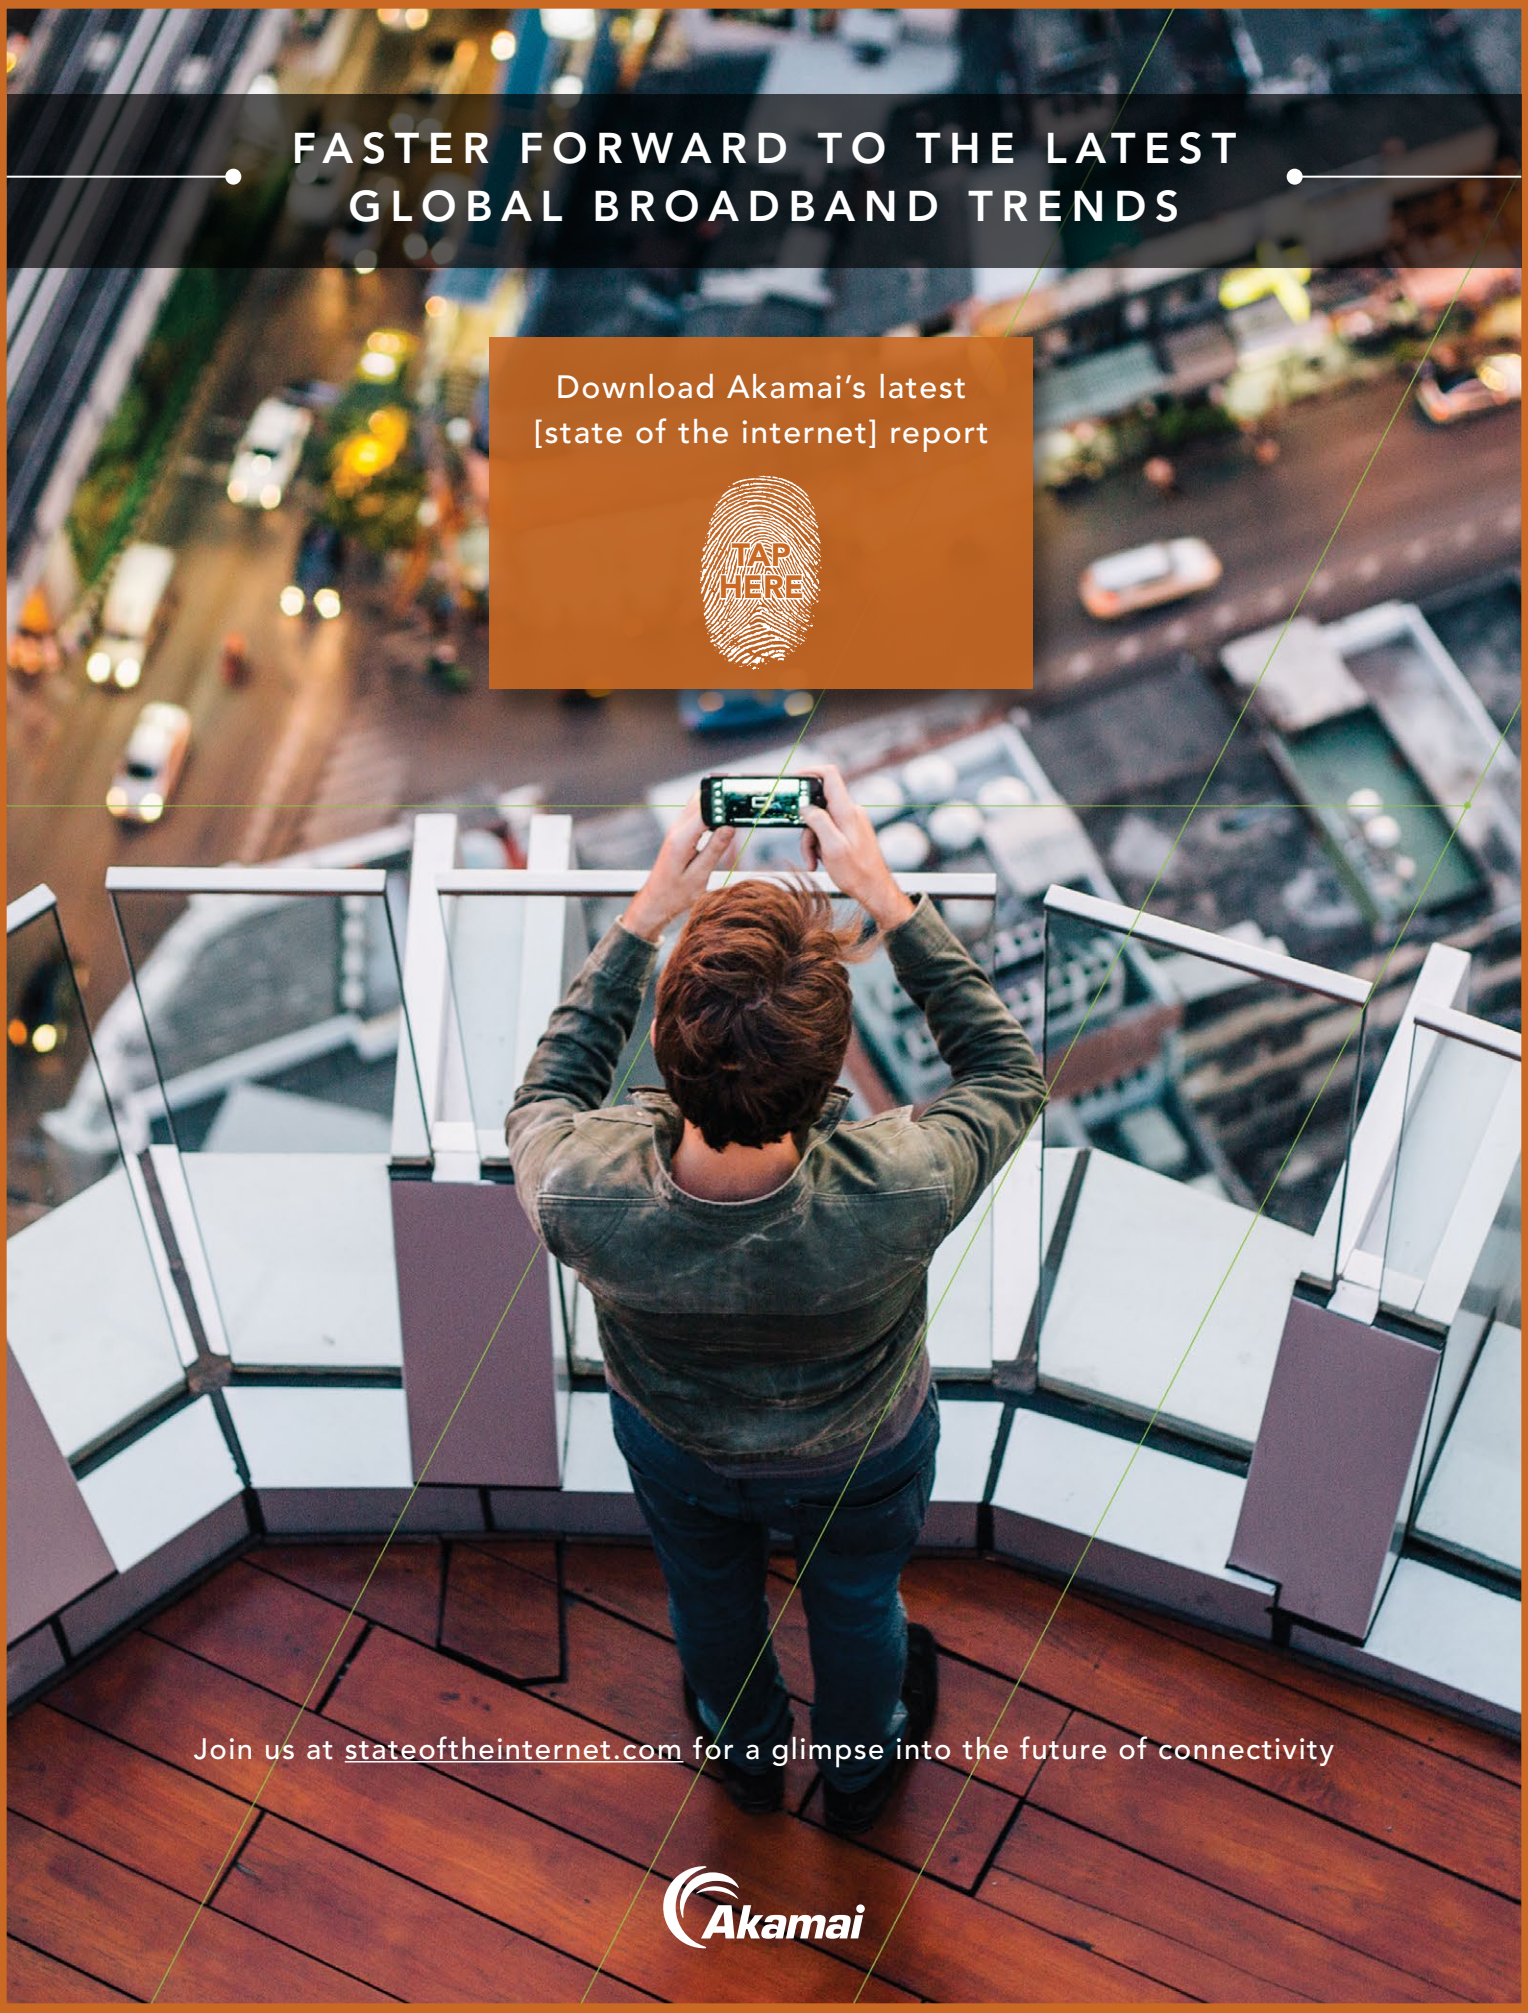
FASTER FORWARD TO THE LATEST
GLOBAL BROADBAND TRENDS
Download Akamai's latest
[state of the internet] report
TAP
HERE
Join us at stateoftheinternet.com for a glimpse into the future of connectivity
Akamai

LETTER FROM THE EDITOR / The Q3 2015 *State of the Internet — Security Report* continues to build on the enhancements we made to the report earlier this year.

Each of our technology platforms collects a distinct data set that reveals a unique view of the Internet. This breadth allows Akamai to compare and contrast the different indicators of attack activity.

We explore which industries among our customer base suffered the highest volume of attacks, which attack techniques and vectors were most common, and where the attack traffic originated.

This report is the first since Akamai launched its Security Intelligence Response Team (***Akamai SIRT***). Combining the teams responsible for intelligence, research, and customer security incident response on the Akamai platforms brings together a cohesive group that sees a large part of the Internet's traffic. The new team's efforts add a greater breadth and depth to the report you are about to read.

We hope you find it valuable.

As always, if you have comments, questions, or suggestions regarding the *State of the Internet — Security Report*, the website, or the mobile applications, connect with us via email at ***stateoftheinternet-security@akamai.com*** or on Twitter at ***@State_Internet***.

You can also interact with us in the *State of the Internet* subspace on the Akamai Community at ***https://community.akamai.com***.

— Akamai Technologies

AT A GLANCE

DDoS attacks compared with Q3 2014

179.66% increase in total DDoS attacks
25.74% increase in application layer (Layer 7) DDoS attacks
198.1% increase in infrastructure layer (Layer 3 & 4) DDoS attacks
15.65% decrease in average attack duration: 18.86 vs. 22.36 hours
52.94% decrease in attacks > 100 Gbps: 8 vs. 17
65.58% decrease in average peak attack bandwidth
88.72% decrease in average peak attack volume
462.44% increase in reflection attacks

DDoS attacks compared with Q2 2015

22.79% increase in total DDoS attacks
42.27% decrease in application layer (Layer 7) DDoS attacks
30.21% increase in infrastructure layer (Layer 3 & 4) DDoS attacks
8.87% decrease in average attack duration: 18.86 vs. 20.64 hours
33.33% decrease in attacks > 100 Gbps: 8 vs. 12
25.13% decrease in average peak attack bandwidth
42.67% decrease in average peak attack volume
40.14% increase in reflection attacks

Web application attacks compared with Q2 2015

96.36% increase in HTTP web application attacks
79.02% decrease in HTTPS web application attacks
21.64% increase in SQLi attacks
204.73% increase in LFI attacks
57.55% increase in RFI attacks
238.98% increase in PHPi attacks

[SECTION][1] ANALYSIS + EMERGING TRENDS

The third quarter of 2015 set a record for the number of distributed denial of service (DDoS) attacks recorded on Akamai's routed network. All told there were 1,510 DDoS attacks — an increase of 180% over what was reported in Q3 2014 and a 23% increase over Q2 2015.

Though the number of DDoS attacks increased, the percentage of attacks that targeted the application layer (Layer 7) dropped 42% in the last quarter. In contrast, application layer DDoS attacks were up 26% compared to Q3 2014.

Infrastructure layer DDoS attacks increased 30% from last quarter and were up 198% over Q3 2014.

While the number of DDoS attacks rose in the last quarter and in the last year, we observed a decrease in average attack duration, as well as average peak bandwidth and volume.

The decreases can be attributed to a few factors. The largest factor is the increasing use of booter-stresser tools. Sites offering booter-stresser tools are purportedly set up to allow administrators to load test their own sites. However, in many cases that is a cover story shrouded in a legal artifact. Many of the sites are simply DDoS-for-hire tools in disguise, relying on the use of reflection attacks to generate their traffic.

Because the vast majority of these sites are subscription-based and usually only allow attacks to last 1,200 – 3,600 seconds (20 – 60 minutes), their use is decreasing the mean length of attacks. In the past, most DDoS attacks were based on infected bots and would last until the attack was mitigated, the malicious actor gave up, or the botnet was taken down. Instead of spending time and effort to build and maintain DDoS botnets, it is far easier for attackers to use booter-stresser tools to exploit network devices and unsecured service protocols.

A review of the data indicates that booter-stresser tools are less capable of the big attacks that infection-based botnets produce. The user login and configuration pages of these tools are almost always hosted behind the protection of a low-cost content delivery network (CDN). This arrangement provides attackers with a layer of perceived anonymity and the ability to launch their attacks, at least for a time, without divulging their point of origin.

In a departure from the last several quarters, Q3 data shows the UK as the top source country for DDoS attacks, responsible for 26% of attacks. China was the second-most prolific source country at 21%, the US came in third (17%), and India and Spain tied for fourth at 7%. Akamai uses application layer traffic as the primary

measurement for the source of DDoS attacks because UDP and other network layer traffic is easily spoofed, disguising the source. Application layer traffic represented a significantly smaller percentage of all DDoS traffic in Q3 than in the past.

Mega DDoS attacks — those measuring 100 Gigabits per second (Gbps) or more — also dropped from 12 last quarter to 8 this quarter. Compared to the Q3 last year, 100+ Gbps-sized attacks have decreased 53%. This is down from the record-setting 17 mega attacks of Q3 2014.

Additionally, in Q3 2015, the largest DDoS attack measured 149 Gbps, a decrease in size from the largest (250 Gbps) last quarter. Of the eight mega attacks, the media and entertainment sector was targeted most frequently, with three attacks.

The online gaming sector was hit particularly hard in Q3 2015, accounting for 50% of the recorded DDoS attacks. Gaming was followed by software and technology, which suffered 25% of all attacks. Internet and telecom was hit by 5% of attacks, a drop from 13% last quarter.

This quarter, the vast majority of web application attacks — 88% — came over HTTP. The remaining 12% came over HTTPS. This drop is dramatic in percentage of HTTPS-based attacks compared to Q2, when Shellshock was used prolifically. Although HTTPS-based web application attacks represent only a small portion of all the web application attacks we observe, they still account for millions of attack alerts each quarter.

This quarter, local file inclusion (LFI) and SQL injection (SQLi) attacks were by far the most prevalent web application attack vectors. The retail industry was hit hardest, receiving 55% of web application attacks, with the financial services industry a distant second, receiving 15% of attacks.

While DDoS attacks frequently relied on booter-stresser sites, web application attacks were more likely to be based on botnets that take advantage of unsecured home-based routers and devices.

For the first time, the security report also includes attack activity observed across the Akamai Edge Firewall, our global platform perimeter. For this report we analyzed the frequency UDP reflection attacks, as well as the corresponding top source ASNs for DDoS reflectors.

In Q3 2015, Akamai tracked several new attack techniques, vulnerabilities and criminal operation campaigns that warranted the release of threat advisories and case studies. These are profiled in more detail in ***Section 5*** and include:

- ***New DDoS reflection techniques***
- ***XOR DDoS***, a Trojan malware that attackers are using to hijack Linux machines
- More attack activity from ***DD4BC*** and the rise of the ***Armada Collective***
- ***Cloudpiercer***, a tool designed to locate the IP addresses of origin servers
- ***A CDN vulnerability*** unveiled at Black Hat USA 2015
- A vulnerability addressed in ***OpenSSL versions 1.0.2d and 1.0.1p***

[SECTION]^2
DDoS ACTIVITY

Despite setting a new record for the number of DDoS attacks in Q3 2015 — a 180% increase over Q3 2014 and a 23% increase over Q2 2015 — we saw decreases in average peak bandwidth and volume, as well as in average attack duration.

The average duration of attacks this quarter was 18.86 hours, a slight drop from 20.64 hours last quarter and a further drop from 22.36 hours in Q3 2014. Average peak bandwidth was 5.15 Gbps, down 25% from last quarter, and average peak volume was 1.57 million packets per second (Mpps), down 43%. Compared to the same period a year ago, peak bandwidth dropped 66% and volume dropped 89%.

The decrease in bandwidth, volume and duration can be attributed to a few factors. One is that the booter-stresser tools used to launch attacks cost money and limit the attacker to a set length of time. In the past, more attacks were based on botnets developed through infection, and attacks could last until they were mitigated, attackers gave up, or the botnet was taken down. Additionally, the booter-stresser tools, which use reflection attack techniques instead of directly generating their own payloads, seem to be less capable of big attacks than botnets.

While the number of application layer (Layer 7) DDoS attacks dropped 42% over last quarter, they were up 26% compared to Q3 2014. Infrastructure layer (layers 3 and 4) DDoS attacks increased 30% from last quarter and were up 198% over Q3 2014.

Attacks measuring 100 Gbps in size or greater also dropped from 12 mega attacks last quarter to 8 mega attacks this quarter. Compared to the same quarter last year, DDoS attacks peaking at 100 Gbps or more decreased 53%.

2.1 / Mega Attacks / In Q3 2015, 8 DDoS attacks registered more than 100 Gbps, as shown in Figure 2-1. This number is down from Q2 2015, when there were 12 mega attacks, and still more of a drop from the record-setting 17 mega attacks of Q3 2014.

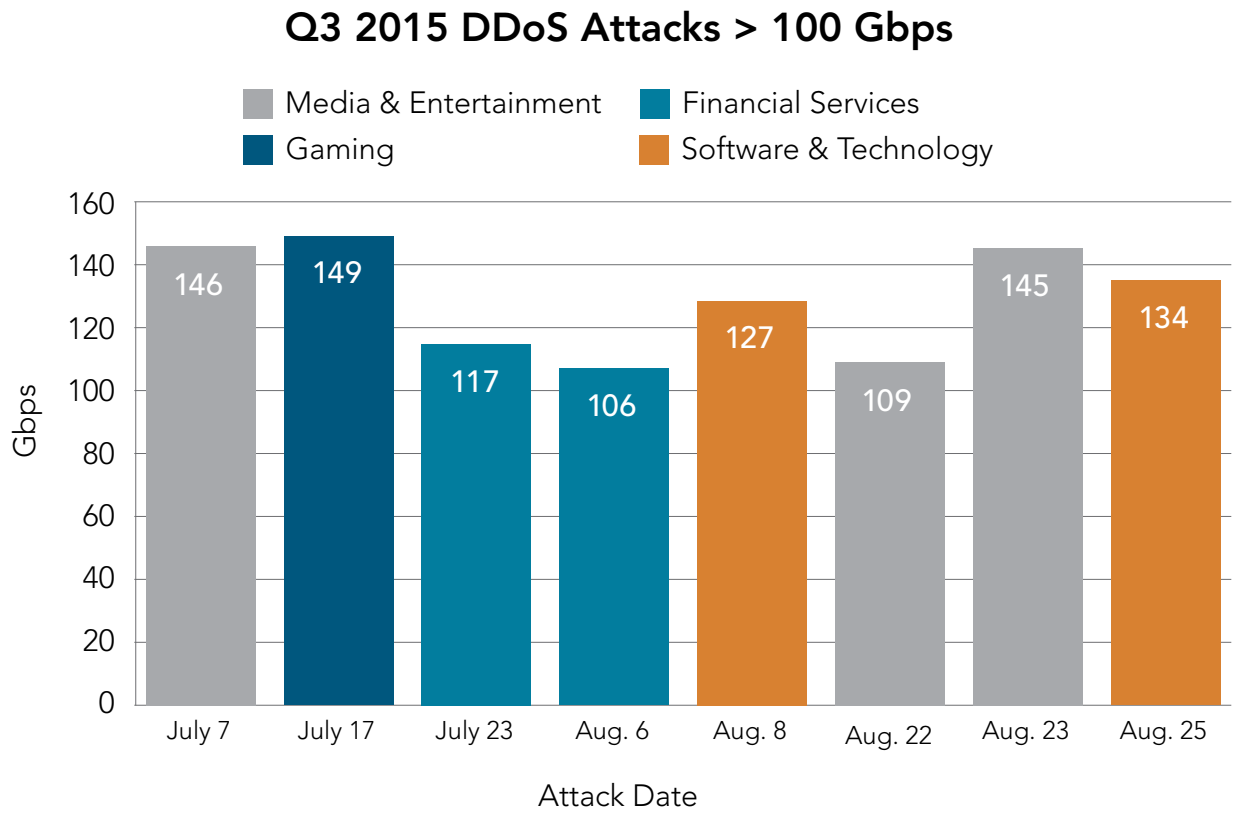

Figure 2-1: Eight DDoS attacks registered more than 100 Gbps in Q3 2015, with the largest registering 149 Gbps

In Q3 2015, the largest DDoS attack measured 149 Gbps, a decrease in size from the largest (250 Gbps) last quarter. Of the eight mega attacks, the media and entertainment sector received the largest share, followed by financial services and software and technology.

The smaller number of large DDoS attacks, along with the lack of any attacks greater than 150 Gbps, is a large part of why the average attack peak bandwidth fell so drastically (25%) in the third quarter. We also saw a large reduction in the number of packets in the larger attacks.

There were five DDoS attacks in Q3 that delivered more than 30 Mpps and only one attack peaked at more than 50 Mpps, although that attack registered an extremely large 222 Mpps, as shown in Figure 2-2. In contrast, there were 18 attacks of 30+ Mpps in Q2 2015. The rate of packets received affects some routers and networks more than the number of bytes because packets require more memory to track, tying up valuable resources.

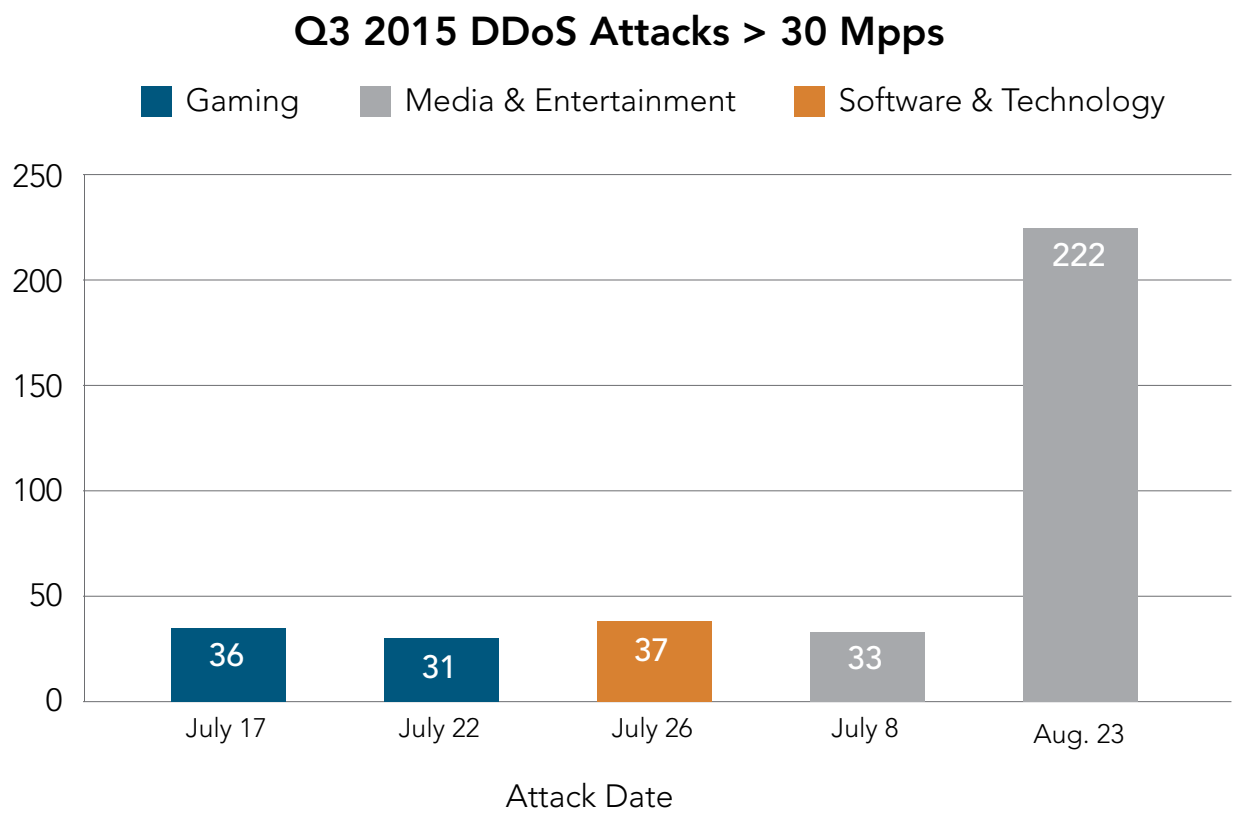

Figure 2-2: Five DDoS attacks delivered more than 30 Mpps in Q3 2015, with one attack much larger than the rest: 222 Mpps

2.2 / DDoS Attack Spotlight: A 222 Mpps Attack / When it comes to DDoS attacks, high bandwidth is usually the most notable metric, but a high packet per second rate can be just as devastating. The attack featured in this spotlight would require at least 145 Gbps of available bandwidth just to withstand the brunt of the attack. Further, the attack generated a record-setting 222 Mpps. The Gbps and Mpps of the malicious traffic were measured at each of six Akamai DDoS scrubbing centers, as shown in Figure 2-3.

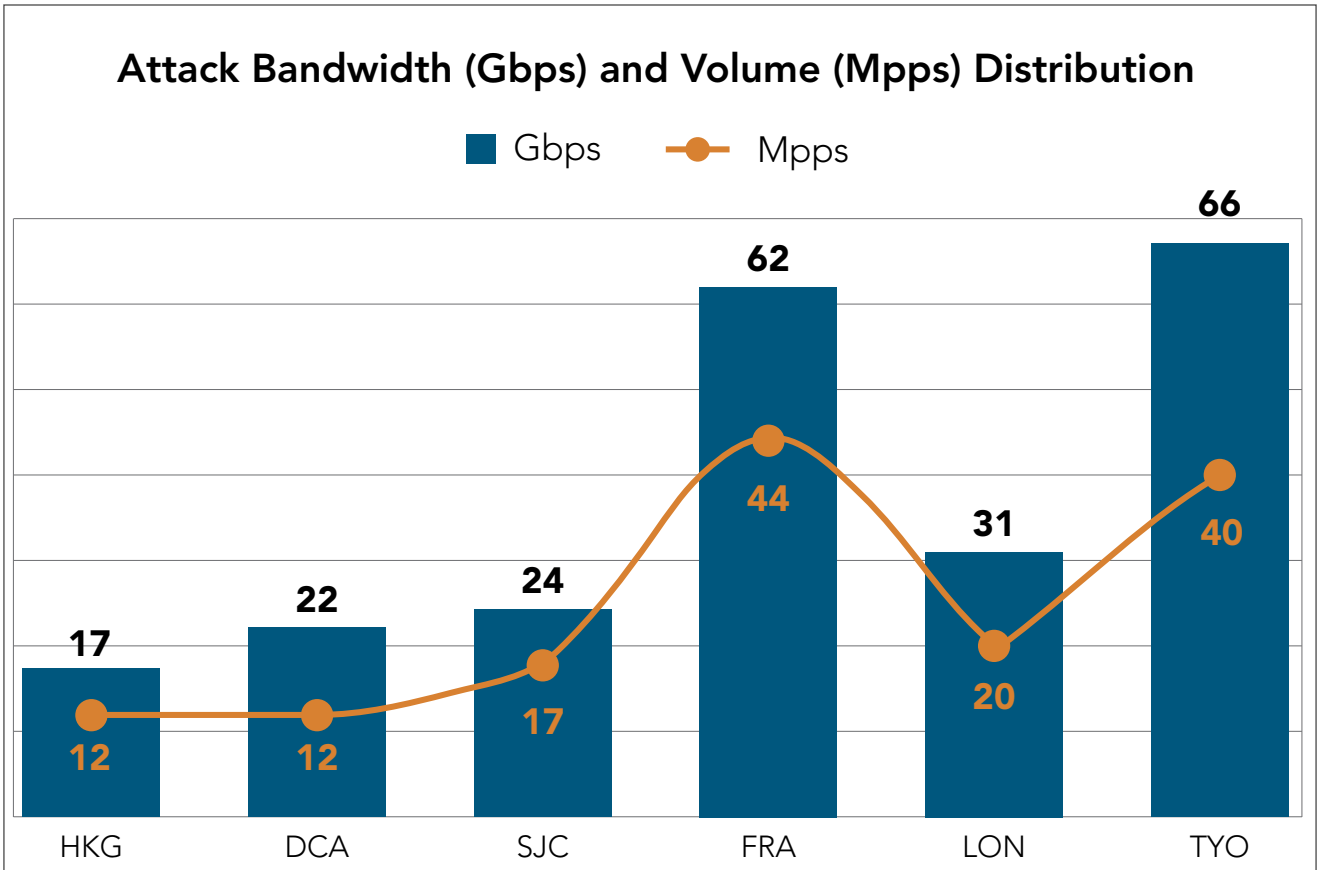

Figure 2-3: Bandwidth (Gbps) and volume (Mpps) of a DDoS attack as measured from Akamai DDoS scrubbing centers in Hong Kong, Washington DC, San Jose, Frankfurt, London and Tokyo

When an attack this large targets a web server, multiple devices are usually traversed before the traffic reaches the target server. The question then becomes, can devices in the path — the edge router, a DDoS mitigation device, a proxy, and any other devices — handle 222 Mpps of malicious traffic and still deliver the load of legitimate traffic?

This DDoS attack generated a high peak traffic using only a single attack vector, a SYN flood. Attacks this large usually employ a combination of at least two attack vectors. Generating such a large load using a single vector may provide a clue as to the size of the botnet.

Payloads / The signatures in Figure 2-4 show the two payload variations observed in this SYN flood attack. One payload has a length of zero, and the other payload has an extra 8 bytes of data.

```
21:28:09.101512 IP X.X.X.X.3478 > X.X.X.X.80: Flags [S], seq 8420, win 21012, options
[mss 729,nop,wscale 8,nop,nop,sackOK], length 0
21:28:09.101517 IP X.X.X.X.4041 > X.X.X.X.80: Flags [S], seq 1612447744:1612447752, win
59258, options [mss 19970,nop,eol], length 8
```

Figure 2-4: Two payload variations used in the SYN flood DDoS attack

Initial packet analysis / Although the malware used in this attack has yet to be analyzed by Akamai SIRT, some initial observations can be made based on packet analysis. For one, there seems to be an anomaly occurring during the creation of these TCP packets. Specifically, some of the TCP options are applied outside of the TCP header and fall into the data portion. In a recent ***analysis of the XOR DDoS*** malware, errors were also found in the creation of the TCP packet related to the calculation of header size.

As shown in Figure 2-5, the first packet contains additional options — the two no operations (NOPs) and the selective acknowledgment (SACK) permitted options. The hex values for those options appear in the data portion of the second packet. Whether the errors in the packets are intentional or not does not diminish the effect of a 222 Mpps flood. This attack does not match what we have previously observed from booter-stresser sites or SYN flood attack scripts. Akamai SIRT will continue to monitor these attacks and provide additional information as it becomes available.

```
SYN flood packet with options applied within header
17:32:00.586714 IP (tos 0x0, ttl 142, id 13523, offset 0, flags [DF], pro-
to TCP (6), length 52)
x.x.x.x.22103 > y.y.y.y.80: Flags [S], cksum 0xa8a4 (correct), seq
4079091712, win 27963, options [mss 1460,nop,wscale 8,nop,nop,sackOK],
length 0
0x0000: 4500 0034 34d3 4000 8e06 a88a XXXX XXXX E..44.@.........
0x0010: XXXX XXXX 5657 0050 f322 0000 0000 0000 .......P.”......
0x0020: 8002 6d3b a8a4 0000 0204 05b4 0103 0308 ..m;............
0x0030: 0101 0402 ....
SYN flood packet with some options outside of the header(NOP NOP 0101 and
SACKok 0402)
17:32:00.586770 IP (tos 0x0, ttl 112, id 24940, offset 0, flags [DF], pro-
to TCP (6), length 60)
x.x.x.x.9309 > y.y.y.y.80: Flags [S], cksum 0x4e52 (correct), seq
4178903040:4178903048, win 4183, options [mss 19202,nop,eol], length 8
0x0000: 4500 003c 616c 4000 7006 03fa XXXX XXXX E..<al@.p.......
0x0010: XXXX XXXX 245d 0050 f915 0000 0000 0000 .......P........
0x0020: 8002 1057 4e52 0000 0204 4b02 0100 0000 ...WNR....K.....
0x0030: 0303 0800 0101 0000 0402 0000 ............
```

Figure 2-5: Errors were made in the packet creation: options are found outside the TCP header

2.3 / DDoS Attack Vectors / As shown in Figure 2-6, network layer attacks continued to account for roughly 95% of all DDoS attack activity.

More than 17% of DDoS attacks utilized the UDP fragment attack vector in Q3 2015, which is a change from last quarter, when SYN floods represented the most popular infrastructure-based attack, as shown in Figure 2-7. Attackers' use of SSDP floods represented 15% of all DDoS activity in Q3, nearly the same as last quarter, when it

represented just under 16% of all DDoS attacks. By comparison, SSDP was the top DDoS attack vector in Q1 2015 and Q4 2014. NTP attacks accounted for 13% of all activity in Q3 and SYN floods decreased to only 12%.

UDP fragments are at least partially a byproduct of other DDoS traffic, such as CHARGEN, DNS, RPC, SNMP and some UDP floods. So while UDP fragments are the largest portion of network layer attack traffic, it is difficult to associate the fragments with their originating type. As a result, the fact that SSDP traffic had previously been the leader for DDoS attack traffic is all the more remarkable.

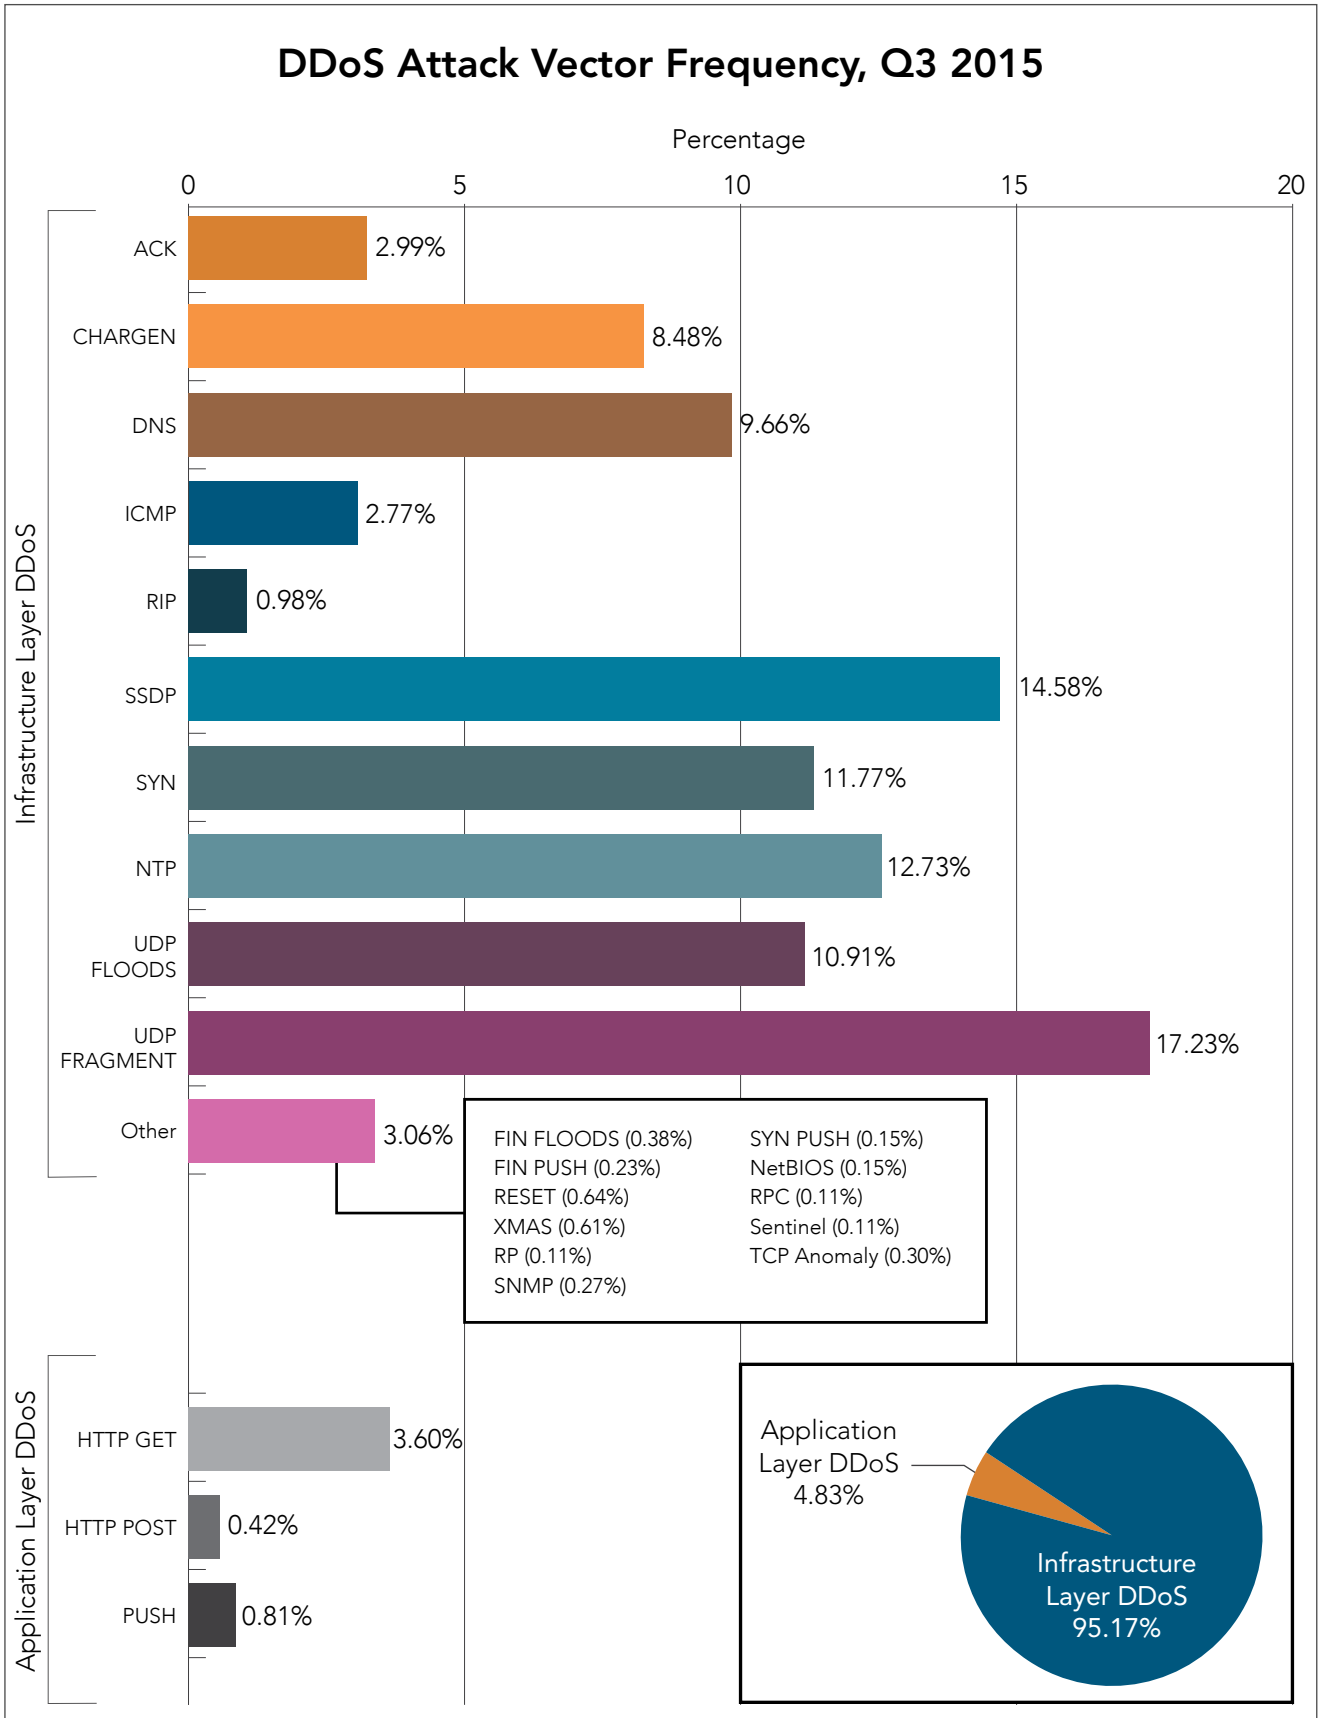

Figure 2-6: Application-layer DDoS attacks were less than 5% of all attacks. UDP fragment floods, SSDP, NTP SYN, UDP floods, DNS, and CHARGEN DDoS attacks all occurred regularly

Although we tracked two dozen attack vectors in Q3 2014, the top 10 vectors were responsible for 95% of the attacks. To better understand the cyclical nature of attacks, we analyzed this subset of attack vectors over the past five quarters. Figure 2-7 shows the frequency of these attack vectors within that subset.

For example, the reduction of SSDP attack traffic and the re-emergence of UDP fragment attacks as the primary tool reflects the cyclical nature of attack tools and methods in the DDoS world. We saw a rapid increase in tools that used SSDP reflection over the last year, as understanding spread of how easily the protocol could be abused. Similarly we saw an increase in NTP attacks in 2014, which will likely recur at the end of 2015 and the beginning of 2016 as new vulnerabilities ***have been recently disclosed*** in the venerable (and vulnerable) Network Time Protocol. That said, not all NTP vulnerabilities produce results that can be used for denial of service attacks. So far the only method being abused is the monlist GET method in NTP queries and few NTP servers appear to still have this vulnerability.

This trend of mostly infrastructure attacks has continued for more than a year, as attackers have relied more and more on reflection attack vectors. Not only do reflection attacks obscure the true IP addresses of the attacker, they also require fewer resources relative to the size of the attack.

That said, DDoS attack scripts for application layer DDoS attacks have been shifting towards the use of non-botnet based resources, such as open proxies on the Internet. This trend, along with the continued abuse of WordPress and Joomla-based websites as GET flood sources, may pave the way to an increase in application-based reflection DDoS attacks that abuse web application frameworks.

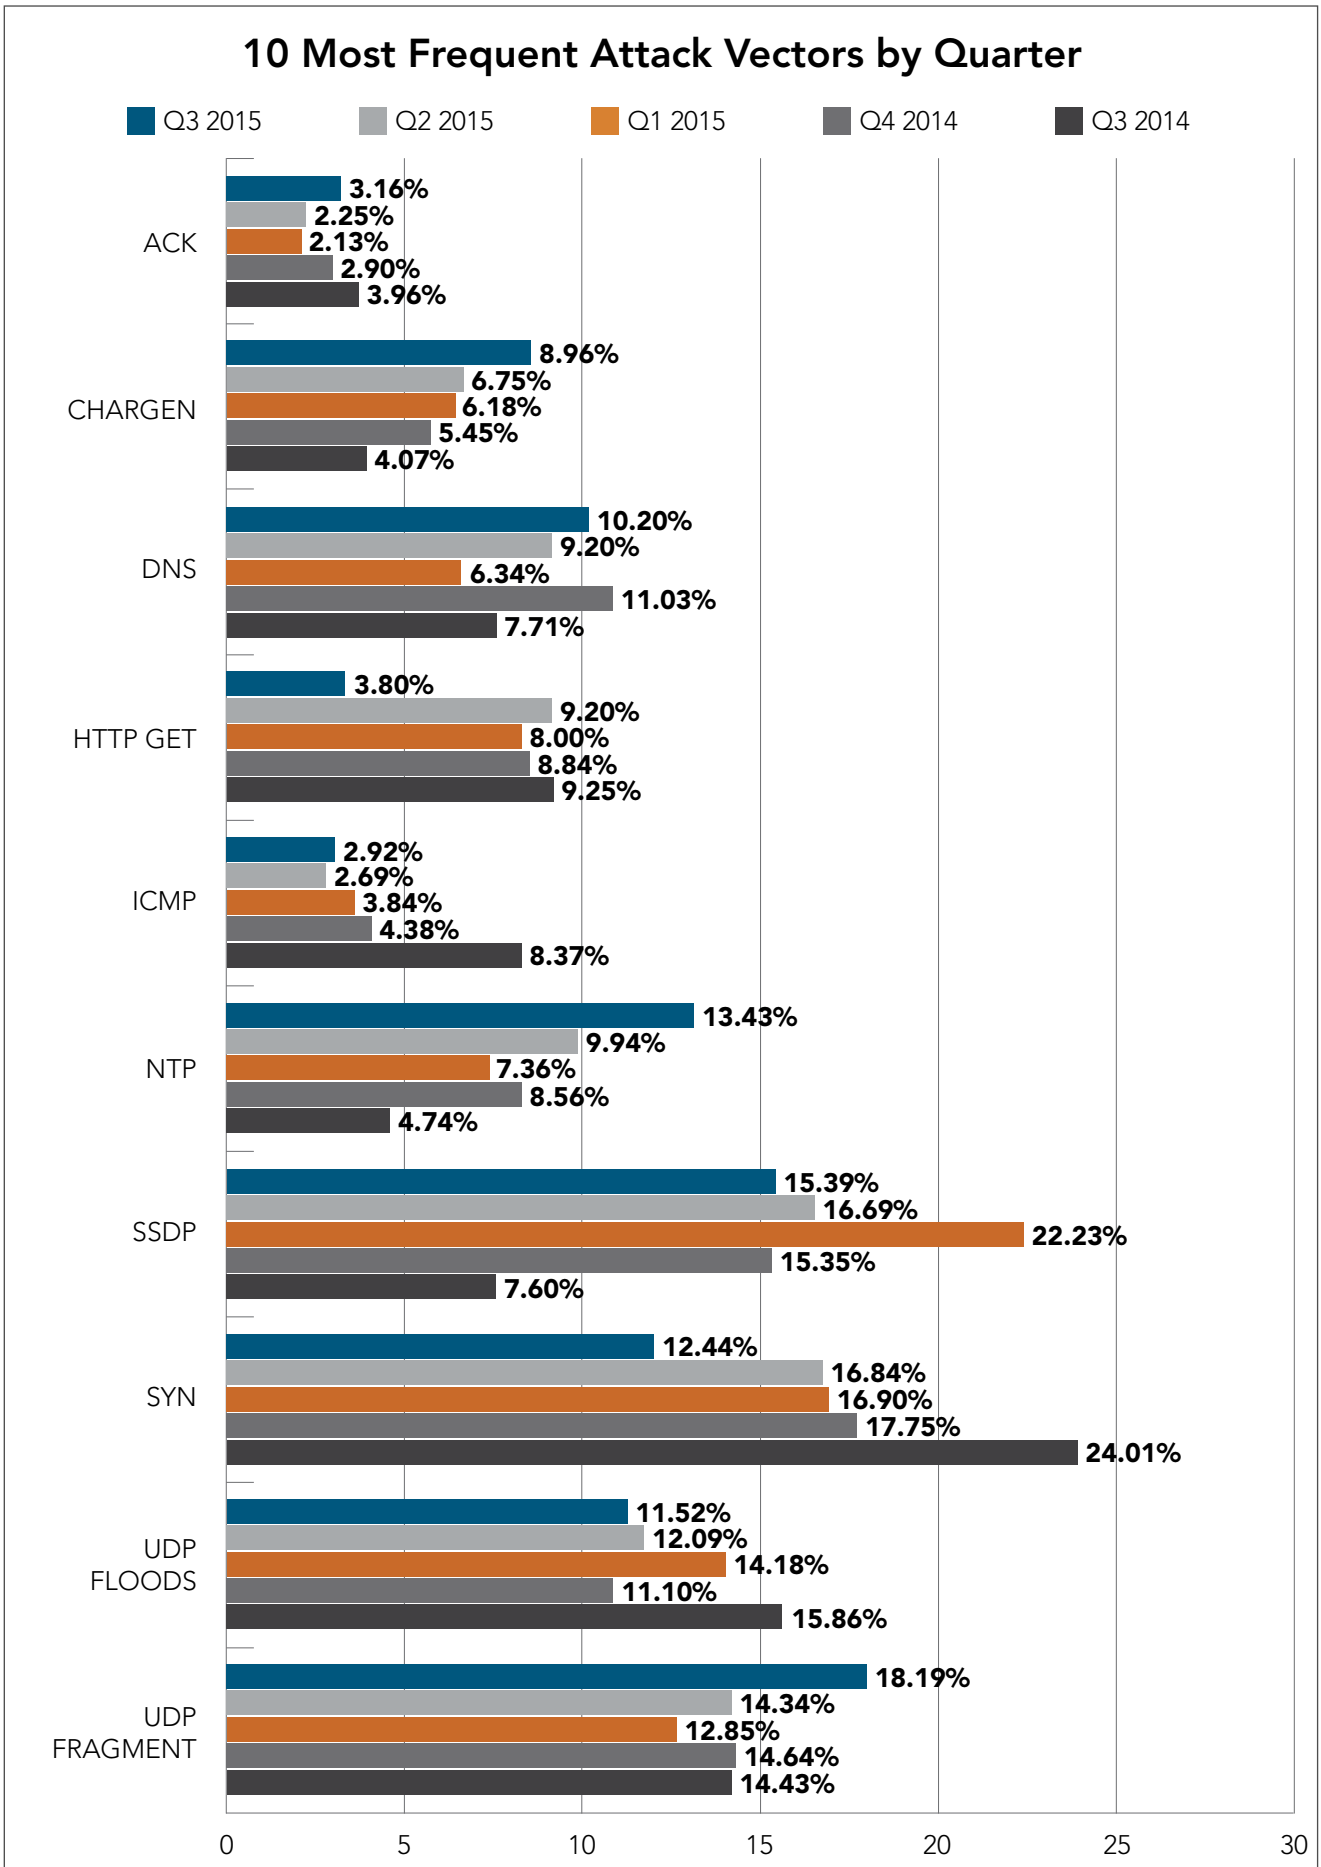

Figure 2-7: Frequency of the 10 most common DDoS attack vectors over the past five quarters

2.4 / Top 10 Source Countries for DDoS Attacks / In a departure from the last several quarters, the UK was the top source country for DDoS attacks in Q3 2015, at 26%, as shown in Figure 2-8. China was the second most prolific source country at 21%, while the US came in third (17%), and India and Spain tied for fourth at 7%.

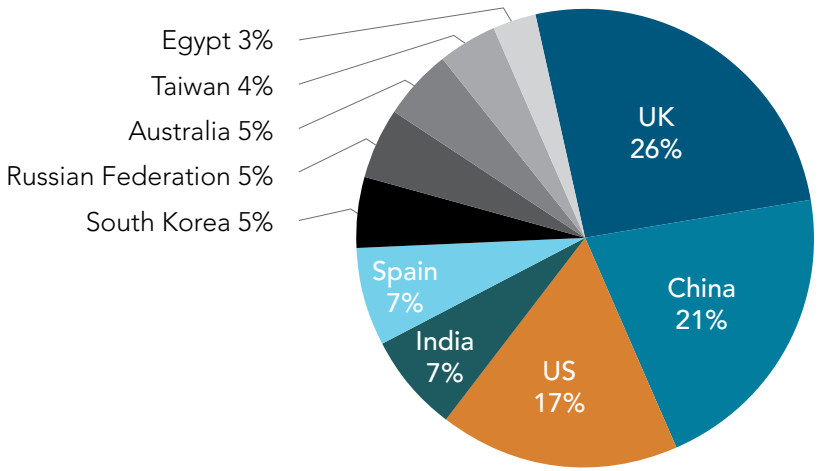

Figure 2-8: The 10 countries sourcing the most DDoS attack traffic in Q3 2015

Figure 2-9 shows the source country results from last quarter and from the prior year. Last quarter China topped the list at 37%, followed by the US (18%), UK (10%) and India (7%). In Q3 2014, the US was the top source country at 24%, followed by China (20%), Brazil (17.5%), and Mexico (14%).

It is important to note that source country is based primarily on application traffic that requires a complete connection. Infrastructure traffic, such as UDP, is easily spoofed, and therefore is not used in this metric. Application layer DDoS attacks, which represent non-spoofed IP addresses, are less prevalent than in the past. In Q2, they represented 10.23% of all DDoS attacks, but in Q3 they only represented 4.83% of all DDoS attacks recorded across the routed platform.

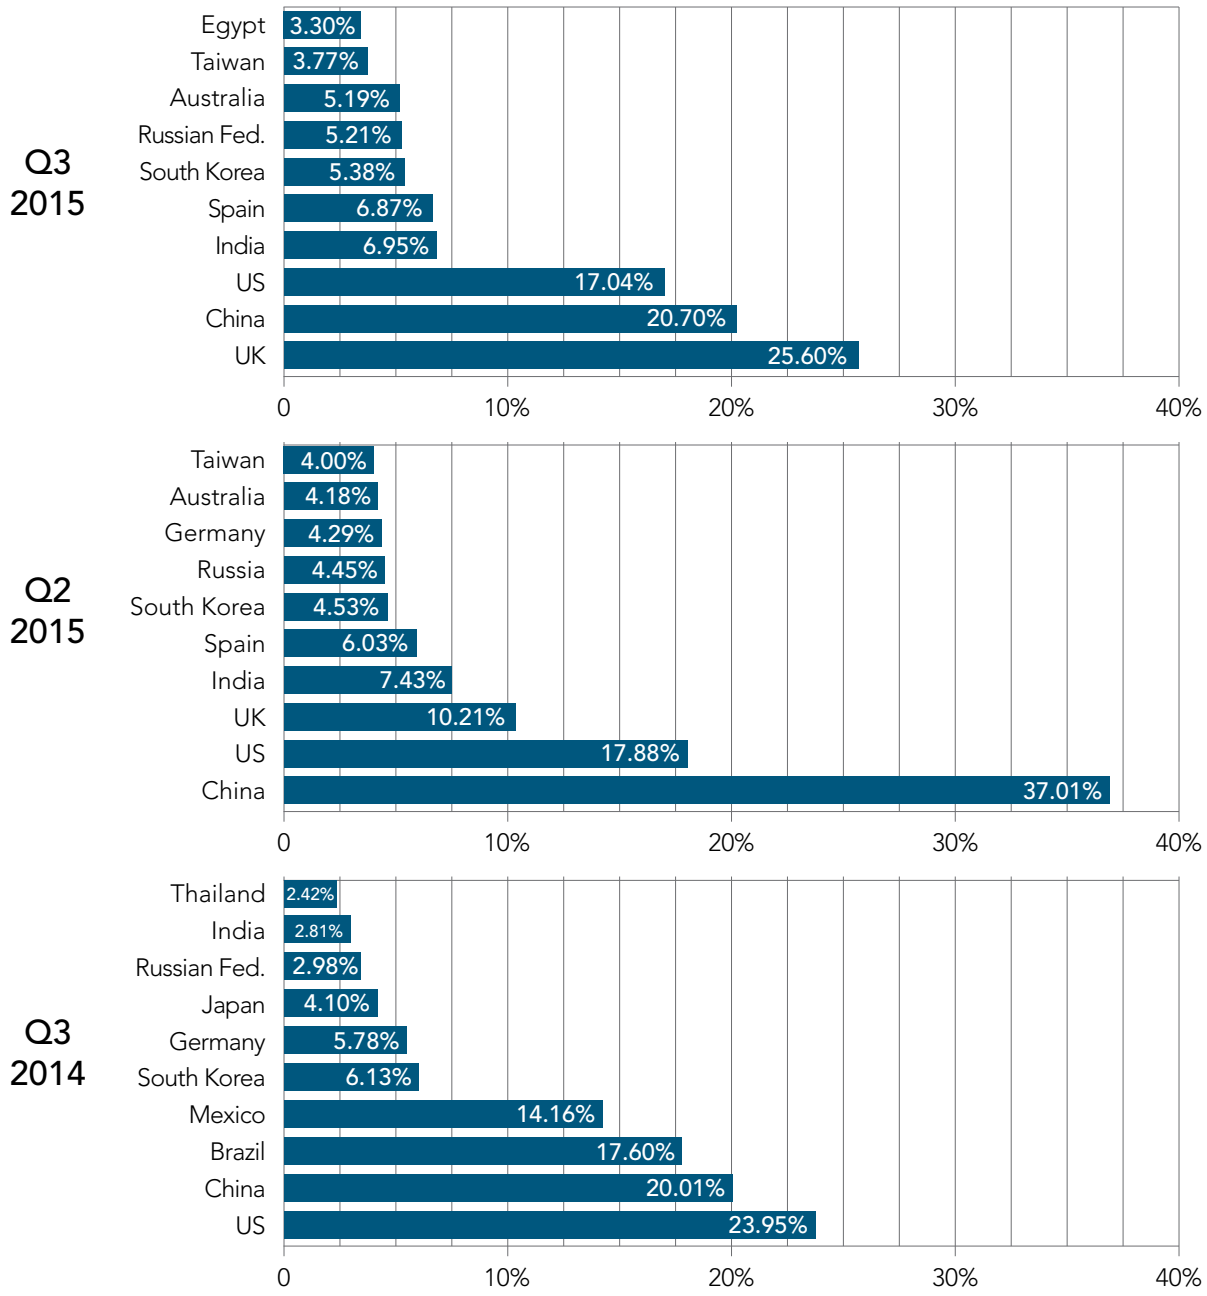

Figure 2-9: The US and China typically are among the top three non-spoofed sources for attacking IP addresses

Since there were fewer non-spoofed attack sources this quarter, the UK's climb to first place is mostly an indication of fewer confirmed attack sources from the other countries, as well as the increase in the number of attacks overall. In fact, the total increase in sources from this quarter to last for the UK is 46.55%, while China and the US dropped 67.32% and 44.31% respectively in total non-spoofed attack sources.

2.5 / DDoS Attacks by Industry / The online gaming sector was hit particularly hard in Q3 2015, accounting for 50% of all DDoS attacks, as shown in Figure 2-10. Gaming was followed by software and technology, which suffered 25% of all attacks in Q3. They were followed by financial services (8%), media and entertainment (5%), Internet and telecom (5%), retail and consumer goods (3%), education (3%) and the public sector (1%).

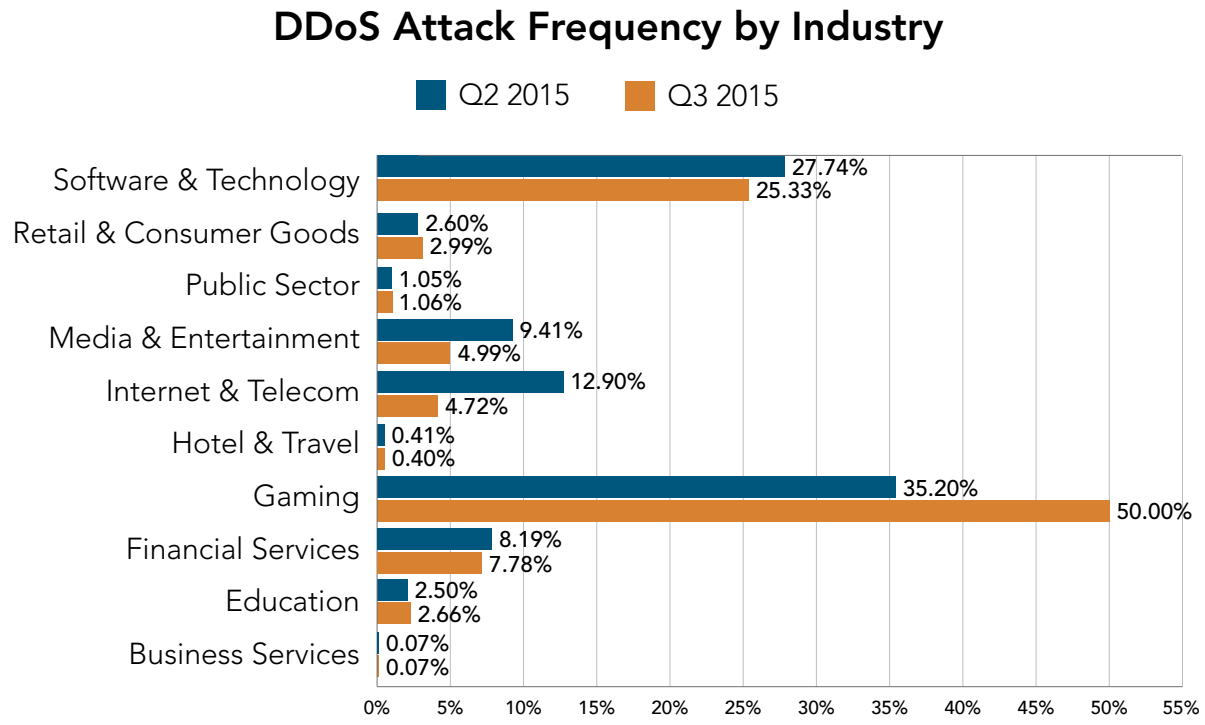

Figure 2-10: Together, online gaming and software and technology were hit by 75% of DDoS attacks in Q3 2015

Online gaming / Online gaming has remained the most targeted industry since Q2 2014. In Q4 2014, attacks were fueled by malicious actors seeking to gain media attention or notoriety from peer groups, to damage reputations and to cause disruptions in gaming services. Some of the largest console gaming networks were openly and extensively attacked in December 2014, when more players were likely to be affected due to the new networked games launched for the holiday season. At the end of 2015, it is likely we'll see a similar pattern emerge again.

Online gaming as a target industry also followed the trend of more reflection DDoS attacks and fewer botnet-based DDoS attacks. This trend was fueled by the availability of booter-stresser sites using reflection attacks and a population of frustrated online gamers, which increases the DDoS risk for this industry.

Software and technology / The software and technology industry includes companies that provide solutions such as Software-as-a-Service (SaaS) and cloud-based technologies. Although this industry saw a slight drop (down from 28% to 25%), relative to other industries last quarter, it actually experienced a slight increase in the number of attacks. The most commonly targeted sub-verticals were chat service providers and non-gaming application developers.

Internet and telecom / The Internet and telecom industry includes companies that offer Internet-related services such as ISPs and DNS providers. It was the target of 5% of attacks in Q3, compared with 13% in the previous quarter. Attackers don't usually target an ISP directly. Instead, the attacks target sites hosted by a provider. The more sites hosted by a provider, the higher the probability that one or more of the sites will be a target for a DDoS attack. The sites can range from personal blogs to commercial sites, and the attackers' motives can vary from politics to extortion.

Financial services / The financial industry includes major financial institutions such as banks, insurance companies, payment providers and trading platforms. The financial industry experienced about the same percentage of all attacks as in Q2 — 8%. Recently, the financial industry has been the focus of various extortion attempts, and the group DD4BC led the way with multiple extortion and DDoS attacks against financial services. As is the case with software and technology, this industry actually saw a slight increase in number of attacks compared to last quarter.

Media and entertainment / The media and entertainment industry had a slight drop in its percentage of attacks, from 9% in Q2 2015 to 5% in Q3 2015. However, a firm in this industry was the target of the largest Mpps DDoS attack recorded to date at 222 Mpps. DDoS attacks on media are usually politically motivated, and attacks can be launched by powerful adversaries, as we saw this quarter.

2.6 / DDoS Attacks — A Two-Year Look Back / When reading about DDoS attacks, most users look at the big attacks and think, "OK, but what should I expect if they're coming after me?" That's the question the graphics below attempt to answer. We're using an interquartile range (IQR) to best represent the data.

Figure 2-11 is a combination of box-and-whiskers chart and a scatter plot of all the DDoS attacks observed by Akamai on the routed platform. We used a logarithmic scale; each interval shows a 10-fold increase in attack size. The scale makes the chart more readable and downplays the outliers — attacks that are exceptionally large or small.

Each dot represents an individual DDoS attack, while the boxes show where the median, the 25th and the 75th percentile fall. For those of us who aren't statisticians, the median is the middle, where 50% of all attacks are larger and 50% are smaller.

While average attack size is highly influenced by the very large and very small attacks, the median hasn't changed much in a year. The median attack size was just under 1.6 Gbps in the third quarter of 2014, just over 1.5 Gbps in Q2, and just under 1.6 Gbps again in Q3 2015. The lower end of the range is slowly increasing as the 25 percentile attacks become larger — 300 megabits per second (Mbps) last year and 425 Mbps in Q3.

Last year at this time, the top of the band was 10 Gbps, while this quarter it was 5.4 Gbps. While the average can be pulled up by a few big attacks, it takes a large number of attacks to change the interquartile range numbers. Attack numbers have increased dramatically over the year with reflection-based attacks leading the way.

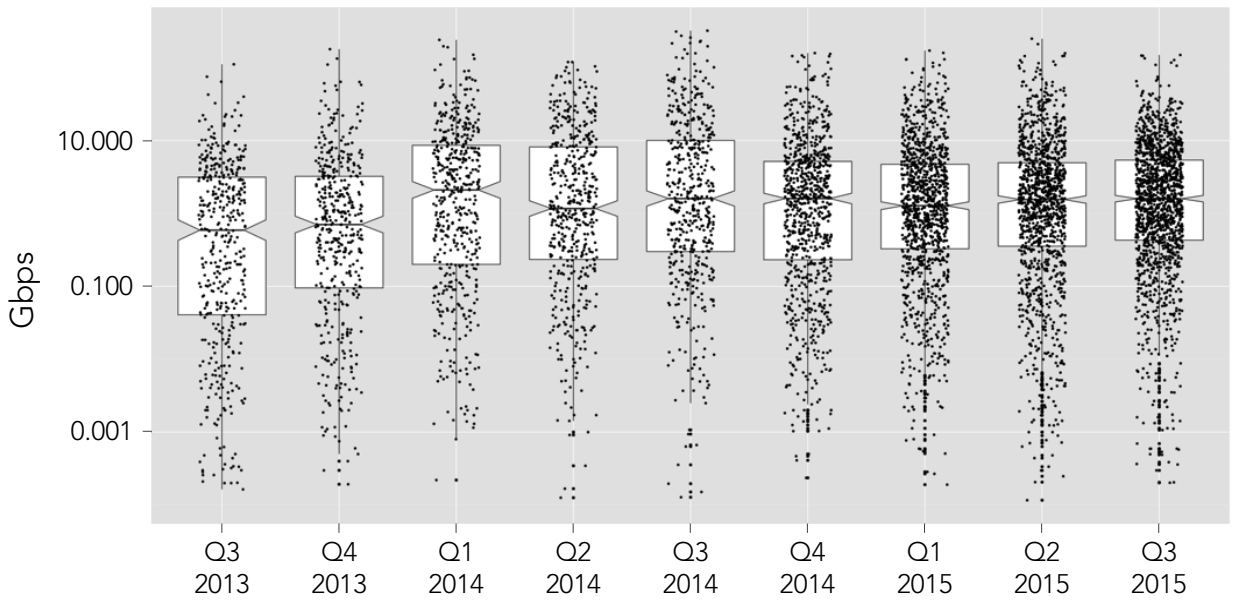

Figure 2-11: DDoS attacks from Q3 2013 to Q3 2015. Each dot represents a DDoS attack, and each interval covers a 10-fold increase in attack size

These attacks don't typically produce the high bandwidth observed in the past from botnet-based attacks. The finite number of reflection sources are likely being used by multiple attackers simultaneously, further reducing the size of each attack.

All of this is a fancy way of saying that if you were to come under DDoS attack, there's an even chance that the DDoS attack size would be somewhere between 400 Mbps and 5 Gbps. This estimate gives you a range to use as you plan your DDoS defenses. Also note that while the attacks falling below 400 Mbps are widely distributed and gradually fall off, there is a large cluster of attacks above 5 Gbps that drops off dramatically above 50 Gbps.

Why we use interquartile range / We're using the IQR to show the median (middle) attack size, as opposed to the mean (average) attack size, because IQR shows the information in a slightly more stable manner, especially as the size of the population of attacks grows.

Before we dive into the shape of the data, here are a few points to know:

- We made a conscious choice to use the median to describe an average attack, rather than the mean. The median is much more resilient to the presence of outliers because it represents the point where half of all attacks are larger and half are smaller.

- The set of observed DDoS attacks include an enormous number of small attacks and a few large ones. For legibility, we chose to use a logarithmic scale, with each interval representing a 10-fold increase.

- There is a notch in each of the boxes centered on the median. The notches show confidence intervals for the median. If the notches for two consecutive boxes overlap, then there is not a statistically significant difference in the median attack size, as is exemplified by Q4 2014 to Q3 2015.

With IQR, population growth tends to create a tighter range in the data. The space between the 75 percentile and 25 percentile marks in the first quarter of 2013 and the same marks in the latest quarter has shrunk considerably. As a result, we start to see the emergence of more statistically significant trends that are less likely to be swayed by a few outliers. This is good; it better showcases the changes in DDoS trends.

Observable patterns / Looking at the time series, a few patterns stand out. First, a significant increase in DDoS attack size occurred in Q1 2014. Prior to that, the first four quarters we tracked (Q1 – Q4 2013) look similar to one another. The upper boundary of the IQR is roughly the same, and three of the four medians are statistically similar.

However, the DDoS trend changed between Q4 2013 and Q1 2014. The upper bound of the IQR increased dramatically (recall, this is a logarithmic scale), as did the median attack size. In Q4 2014, we see another change — a statistically significant drop in the upper bound of the IQR, but the median attack size remained the same. The size of the large DDoS attacks appears to be clumping closer to the median.

2.7 / Reflection DDoS Attacks, Q3 2014 – Q3 2015 / Starting this quarter, we're introducing what is known as a Sankey graphic. Sankey diagrams help to visualize energy, material or cost transfers between processes. This type of diagram was created by Irish Captain Matthew Henry Phineas Riall Sankey, who used such a diagram in 1898 to show the energy efficiency of a steam engine.

The Sankey visualization in Figure 2-12 shows how DDoS reflection attacks have trended in the last calendar year. Through the routed network, we tracked nine infrastructure layer reflection DDoS vectors of which the three newest (and smallest) are ***NetBIOS, Sentinel and RPC.*** Sentinel and RPC were first observed by Akamai this quarter. Every attack vector displayed was involved in a DDoS attack in Q3 2015. The most used vectors correlate with the number of Internet devices that use these specific service protocols for legitimate purposes.

On the left, as indicated by the height of the label, we see that SSDP, NTP, DNS, and CHARGEN were the most used reflection DDoS vectors. As the top vector, SSDP shows a steady increase from Q3 2014 to Q3 2015. The use of the attack peaked in Q1 2015, paused, and then continued an upward trend in Q3 2015.

On the right, we see a steady increase in the use of reflection DDoS as an attack method. The number of reflection DDoS attacks overall has increased dramatically over the last year, and the diagram shows that reflection attacks are a big part of the current landscape.

In a reflection DDoS attack, a malicious actor begins by sending a query to a victim IP address. The victim is an unwitting accomplice in the attack. The victim could be any device on the Internet that exposes a reflectable UDP service. The attacker's query is spoofed to appear to originate from the attacker's target. The attacker uses an automated attack tool to send malicious queries at high rates to a large list of victims, who will in turn respond to the target.

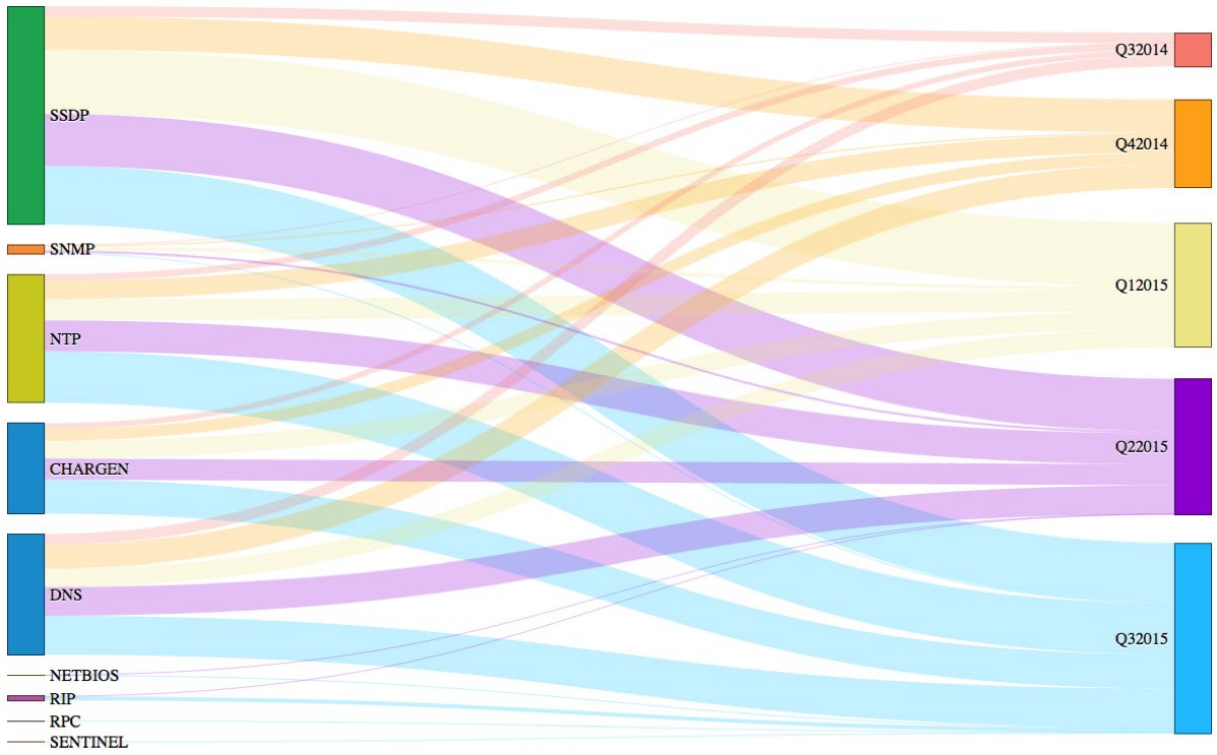

Figure 2-12: This Sankey visualization shows the trends in the types of DDoS reflection attacks used over the five quarters

The percentage of all DDoS attacks that use a reflection attack method is indeed growing. Figure 2-13 shows the percentage of mitigated DDoS attacks that used reflection. In Q3 2014 only 5.90% of all attack campaigns included reflection, but 33% used reflection in the most recent quarter. Keep in mind that in many DDoS attack campaigns multiple attack vectors are utilized simultaneously; one third used at least one reflection vector in Q3.

A big takeaway from the Sankey graph is that malicious actors are finding it more profitable to choose reflection over infection. Instead of spending time and effort to build and maintain DDoS botnets, it is far easier for attackers to exploit network devices and unsecured service protocols. This methodology has been applied to the DDoS-for-hire ecosystem.

Reflection attacks are further facilitated by the connectionless nature of UDP. Unlike TCP, which by virtue of the three-way handshake verifies the source of a request, UDP will always reply to the source IP of a crafted request. This behavior allows for the sending of malicious queries with spoofed source IP addresses. As a result, a flood of replies ends up in the hands of an unfortunate target.

Quarterly Percentage of Reflection-Based DDoS Attacks, Q3 2014–Q3 2015

Quarter	Percentage
Q3 2015	33.19 %
Q2 2015	23.68 %
Q1 2015	21.71 %
Q4 2014	15.51 %
Q3 2014	5.90 %

Figure 2-13: The percentage of DDoS attacks that were based on reflection vectors has increased in each of the past five quarters

[SECTION]3 WEB APPLICATION FIREWALL ACTIVITY

Akamai's research teams concentrated their analysis on nine common web application attack vectors — a cross section of many of the most common categories on industry vulnerability lists. Akamai's goal was not to validate a vulnerability list but to look at some of the characteristics of the attacks as they transit our large network.

As with all sensors, the data sources we use have varying levels of confidence. For this report, we aimed for the lowest rate of false positives and focused on the most highly-used web application attack vectors identified within our threat landscape.

3.1 / Web Application Attack Vectors /

Last quarter, we added two data points to the web application attacks on which we are reporting: xss and Shellshock. Including events based on Shellshock nearly doubled the number of attack events we analyzed in Q2, with 173 million Shellshock attacks against Akamai customers in that one quarter. Shellshock also significantly shifted the balance of attacks over HTTP vs. HTTPS, in large part because these attacks were mostly carried out over HTTPS in Q2 2015. The Shellshock bug was first announced in September 2014 and received heavy media attention. As a result, this bug is likely to be patched on most systems. The number of attempts to exploit it should continue to drop.

The proliferation of botnets built from home router devices is causing an increase in Shellshock attempts over HTTPS as criminals attempt to compromise routers by exploiting default login credentials and unpatched firmware still vulnerable to Shellshock.

WEB APPLICATION ATTACK TYPES

SQLi / SQL injection is an attack where adversary-supplied content is inserted directly into a SQL statement before parsing, rather than being safely conveyed post-parse via a parameterized query.

RFI / Remote file inclusion is an attack where a malicious user abuses the dynamic file include mechanism, which is available in many web frameworks, and loads remote malicious code into the victim web application.

PHPi / PHP injection is an attack where a malicious user is able to inject PHP code from the request itself into a data stream, which gets executed by the PHP interpreter, such as by use of the eval() function.

MFU / Malicious file upload (or unrestricted file upload) is a type of attack where a malicious user uploads unauthorized files to the target application. These potentially malicious files can later be used to gain full control over the system.

CMDi / Command injection is an attack that leverages application vulnerabilities to allow a malicious user to execute arbitrary shell commands on the target system.

(*Continued on next page*)

While botnets are fueling Shellshock attacks, SQLi and LFI attacks remain the dominant attack vectors due to tool availability — attackers frequently use free and open-source tools for SQLi and LFI attacks to find and exploit vulnerabilities in sites.

The third quarter was also notable for an increase in WordPress plugin attack attempts, not only for popular plugins but also for less-known vulnerable plugins.

3.2 / Web Application Attacks Over HTTP vs. HTTPS / This quarter the majority of attacks — 88% — came over HTTP. The remaining 12% came over HTTPS. In the big picture, HTTPS-based attacks represent only a small portion of the attacks we see, yet they account for millions of attack alerts each quarter.

Given that a large percentage of websites either do not use HTTPS for all of their web traffic, or use it only for safeguarding certain sensitive transactions (such as login requests), the comparison between HTTP and HTTPS should be used only to understand attack trends across the two communication channels.

LFI / Local file inclusion is an attack where a malicious user is able to gain unauthorized read access to local files on the web server.

JAVAi / Java injection is an attack where a malicious user injects Java code, such as by abusing the Object Graph Navigation Language (OGNL), a Java expression language. This kind of attack became very popular due to recent flaws in the Java-based Struts framework, which uses OGNL extensively in cookie and query parameter processing.

XSS / Cross-site scripting is an attack that allows a malicious actor to inject client-side code into web pages viewed by others. When an attacker gets a user's browser to execute the code, it will run within the security context (or zone) of the hosting web site. With this level of privilege, the code has the ability to read, modify and transmit any sensitive data accessible by the browser.

Shellshock / Disclosed in September 2014, Shellshock (***CVE-2014-6271***) is a vulnerability in the Bash shell (the default shell for Linux and Mac OS X) that allows for arbitrary command execution by a remote attacker. The vulnerability had existed in Bash since 1989, and the ubiquitous presence of Bash makes the vulnerability a tempting target.

Encrypted connections (over HTTPS do not provide any additional attack protection for web applications. Attackers will likely shift to HTTPS to follow vulnerable applications.

Looking at the Q3 data, we see three possible web application attack trends. First, attacks are coinciding with more sites adopting Transport Layer Security (TLS) HTTPS, as opposed to SSL. Second, attackers are attempting more stealthy attacks over HTTPS, possibly to evade simple intrusion detection systems. And finally, attackers may have fully encrypted connections and are defaulting to HTTPS attacks.

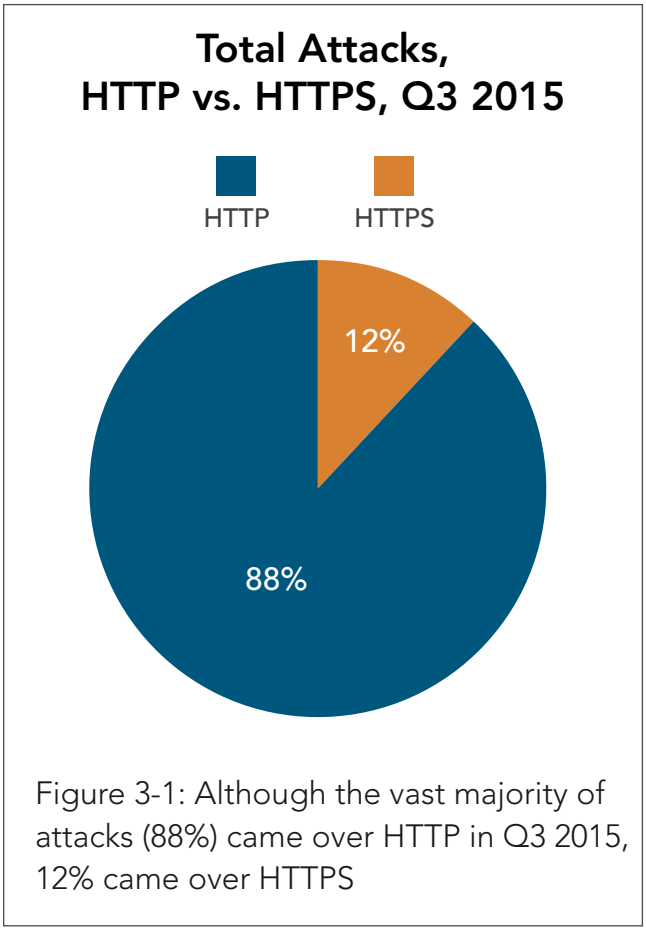

Figure 3-1: Although the vast majority of attacks (88%) came over HTTP in Q3 2015, 12% came over HTTPS

With more Internet sites adopting TLS-enabled traffic as a standard security layer, attackers may follow suit. Or, it could be that attackers aren't looking solely to penetrate a site but to target a back-end database, which is most likely accessed via HTTPS.

Figure 3-2 shows a week-by-week view of the most-used web application attack vectors over HTTPS in Q3 2015. LFI was exploited most often, followed by SQLi. Only in week 30 was SQLi more prevalent.

Figure 3-3 provides a similar view over HTTP. Again LFI and SQLi are exploited most often. SQLi attacks exceeded LFI in three weeks: 25, 26, and 32.

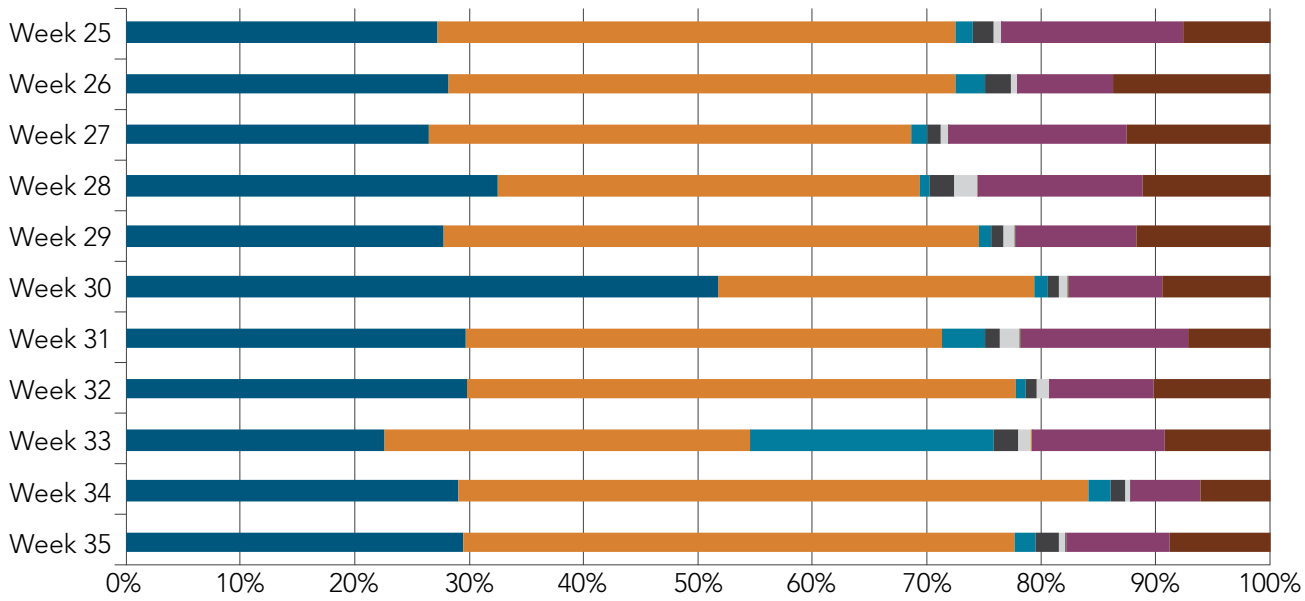

Figure 3-2: The web application attack vectors for attacks over encrypted HTTPS mitigated by Akamai for each week in Q3 2015. LFI was exploited most of the time, followed by SQLi

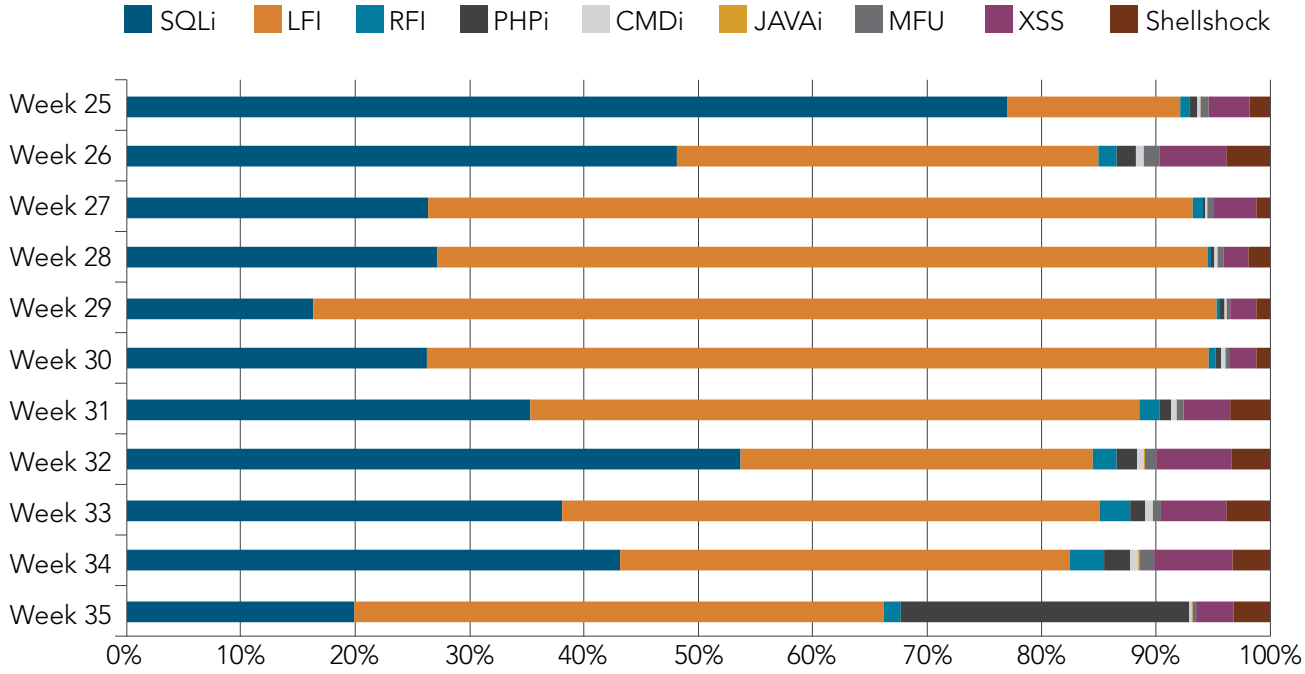

Figure 3-3: The web application attack vectors for attacks over HTTP (not encrypted) mitigated by Akamai for each week in Q3 2015. LFI and SQLi were exploited most of the time

3.3 / Top 10 Source and Target Countries for Web Application Attacks /

In Q3 2015, the US was the main source of web application attacks, accounting for 59% of attack origin traffic, as shown in Figure 3-4. China was the second largest source country at 11%, followed by Brazil (7%), Russia (7%), Bulgaria (4%), Ukraine and the UK (3% each), and the Netherlands, Turkey and Moldova (2% each). Due to the use of tools to mask the actual location, the attacker may not have been located in the country detected. These countries represent the IP addresses for the last hop observed.

The web application attacks we analyzed occurred after a TCP session was established. Therefore, the geographic origins of the attack traffic can be stated with high confidence. Countries with a higher population and higher Internet connectivity are often observed as the source of web application attack traffic.

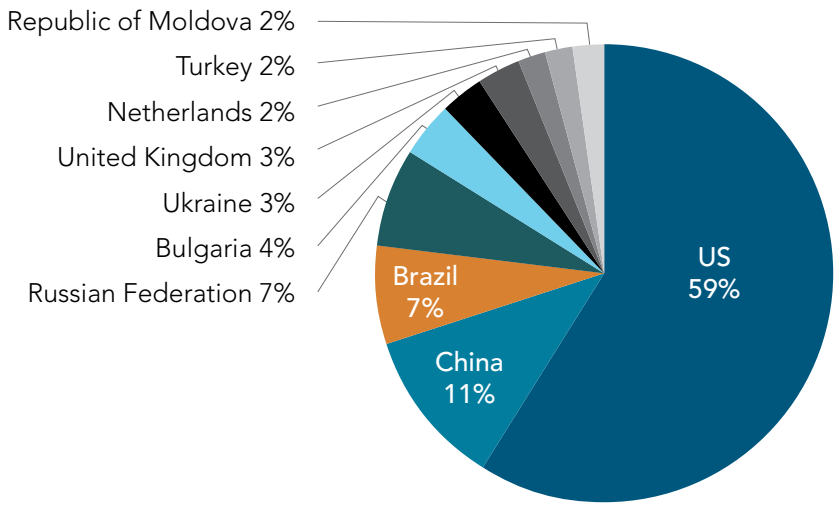

Figure 3-4: Top 10 source countries for web application attacks. The US was the main source, accounting for 59% of attack origin traffic

We can generate significant insight from an analysis of traffic based on the Autonomous System Number (ASN) assigned to traffic in association with Border Gateway Protocol (BGP) routing. The ASN uniquely identifies each network on the Internet with a high degree of reliability. Although an IP address can be spoofed easily, the ASN of the originating traffic is almost always beyond the power of the attacker to change.

In Q3, ASNs also show the US was the top source of malicious web traffic recorded within the Akamai Kona Site Defender infrastructure, followed by China and Russia, as shown in Figure 3-5.

The top three attacking ASNs were associated with a virtual private system (VPS) owned by a well-known cloud provider. While it is easy to set up a system in the cloud, it is hard to secure it. As a result, many of the systems that are set up each day are compromised easily and used in a botnet or other attack platform.

There are three reasons why it is hard to secure a cloud platform. First, anyone can establish a VPS, but it takes a large base of knowledge and motivation to properly configure a system securely. And just like a physical system, one misconfiguration or forgotten patch can make a system insecure. Second, it is easier, cheaper and less traceable to set up malicious servers in the cloud than on owned hardware. Bringing up a system that can be created and torn down in seconds with a few commands is a powerful incentive for legitimate users and attackers alike. Third, while many VPS providers have extensive tools to identify fraud and the theft of system keys, identifying a command and control (C&C, C2) structure for a botnet is much more difficult and might be indistinguishable from normal web traffic.

Country	Type	Total Web Attacks
US	VPS Provider	7,938,425
US	VPS and Colocation Provider	4,121,605
US	VPS Provider	3,882,181
CN	N/A	11,031,342
CN	N/A	8,079,761
CN	N/A	5,989,781
RU	ISP	5,749,782
RU	ISP	4,647,300
RU	ISP	3,730,969

Figure 3-5: Highest sourcing malicious traffic of the top three source countries identified

This quarter, the US had the unpleasant distinction of being both the top source of web application attacks and their top target. Given that many companies have their headquarters and IT infrastructure in the US, this makes sense. Seventy-five percent of web application attacks targeted the US, while only 7% targeted the UK, 6% targeted Brazil and 4% targeted India. Germany and China only found themselves on the receiving end of 2% of all web application attacks, and Australia, Canada, Japan, and Singapore were hit by 1% apiece.

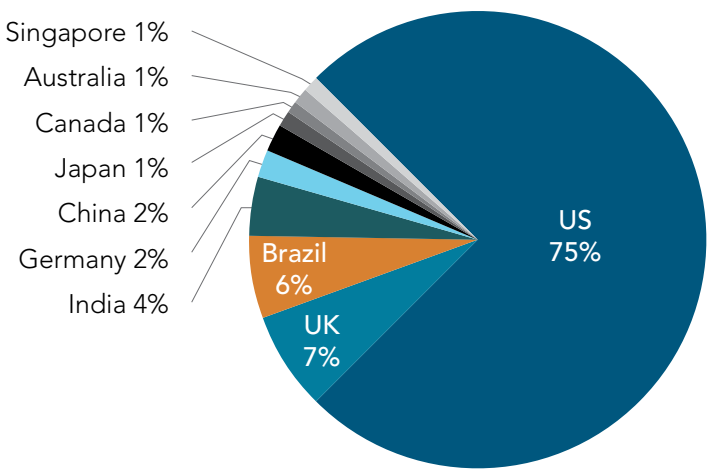

Figure 3-6: The US was targeted by 75% of web application attacks in Q3 2015

3.4 / Web Application Attacks by Industry / This quarter, retail suffered the vast majority of web application attacks — 55% as shown in Figure 3-7. Financial services suffered 15% of attacks, followed by media and entertainment (8%), public sector and hotel and travel (7% each), technology (3%), consumer goods and gaming (2% each), and business services (less than a percent).

Retailers are targeted for DDoS attacks, but they are also targeted for web application layer attacks for significant reasons. Retailers have large amounts of valuable information in their databases, and if an adversary is able to find an SQL injection vulnerability, the attacker can access the retailer's information. Retailers also have a large number of visitors to their websites. As a result, attackers will find and exploit cross-site scripting vulnerabilities to deface retailers' websites, causing a loss of

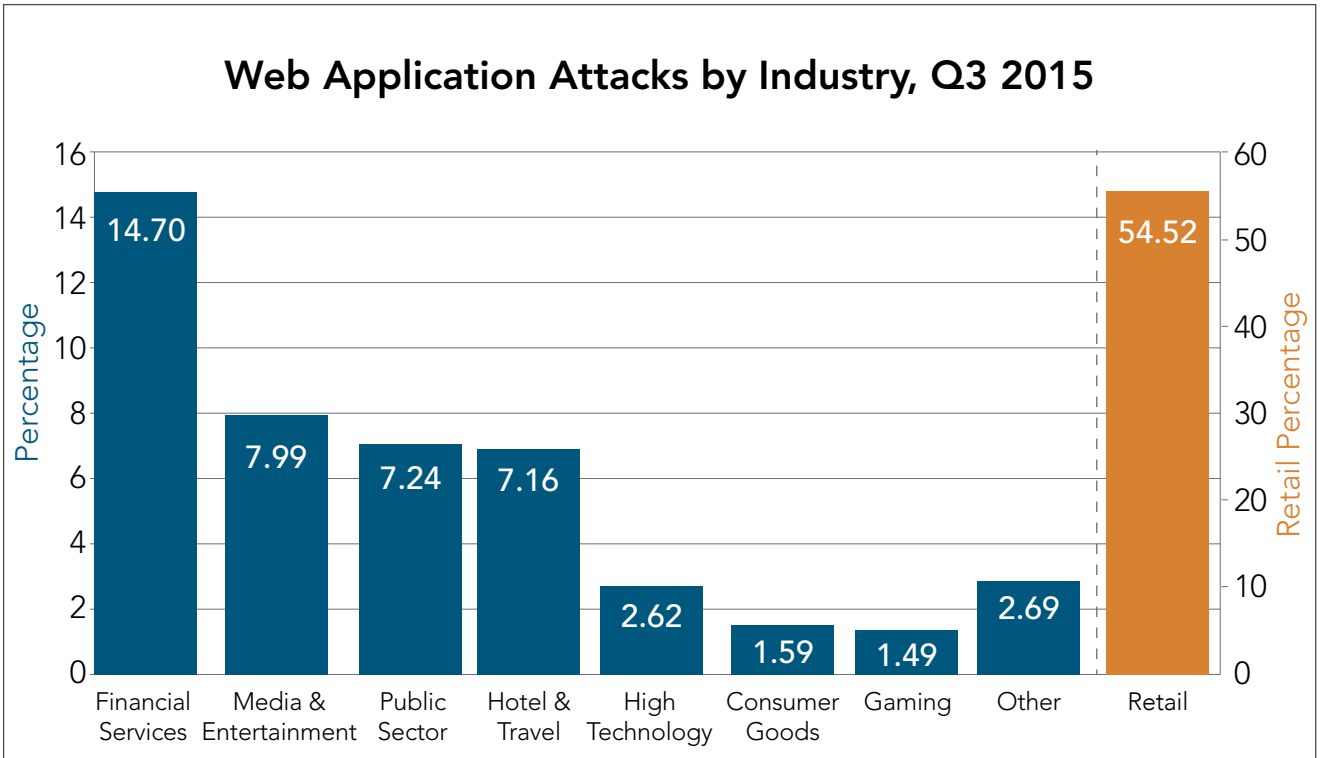

Figure 3-7: Retail suffered the vast majority of web application attacks in Q3 2015—55%

trust among customers. Alternately, the attacker may use a compromised site for a watering hole attack, loading malware on site visitors' computers. Retailers may also be a target for unvalidated requests. For example, if an attacker can control the price of the item being purchased, items may be sold for an amount much different than the retailer intended. Merchants need to be cognizant of all possible vectors through which their web applications may be compromised.

In contrast, network layer DDoS attacks are more prevalent in the gaming sector, primarily because they are an attempt to disrupt the service. Additionally, the gaming industry often has a smaller attack surface for web application attacks than industries such as online retailers.

3.5 / SQLi and LFI Attacks by Target Industry / Much like last quarter, in Q3 2015 the industries subjected to the greatest number of malicious SQLi and LFI requests were the retail and financial services verticals, as shown in Figure 3-8.

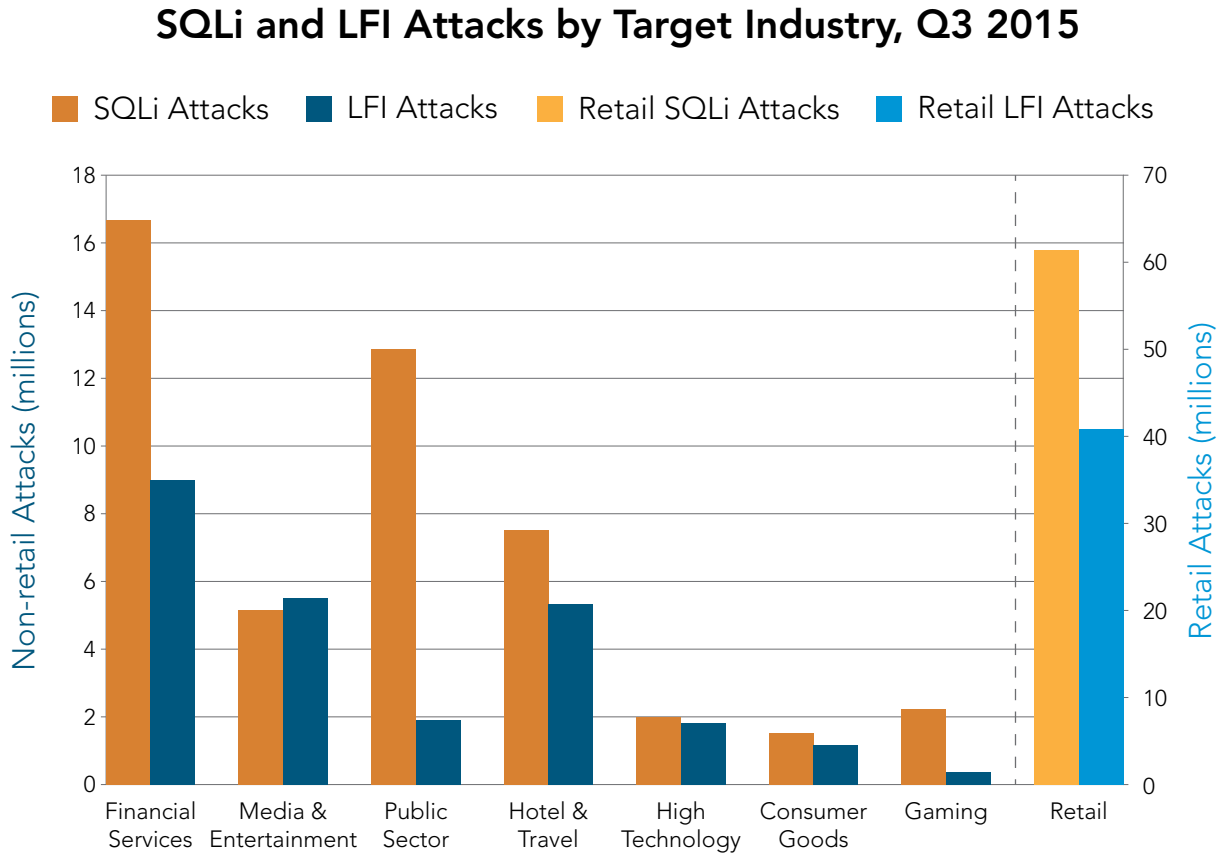

Figure 3-8: Financial services and retail were the industries most targeted by SQLi and LFI attacks

The most common attack vector, SQLi, takes advantage of improper coding of web applications that allows attackers to inject SQL statements, or fragments of SQL statements, into predefined back-end SQL statements such as those used by a login form. This may in turn allow the attacker to gain access to data held within a database or to perform other malicious actions. SQLi and LFI attacks were attempted against Akamai customers more than any other web application attack vector, and the targets of those attackers were most often financial services and retail firms.

LFI attack attempts can be seen in server logs by examining them for indicators of directory traversal attempts. These attempts appear as repeated strings of `../` ending with a filename on a UNIX-based server, or a `..\` on a Windows-based server. The LFI attack will attempt to read sensitive files on the server that were not intended to be available publicly, such as password or configuration information. LFI attacks were the second most common attack vector in Q3 2015, most frequently targeting retail and financial services sites.

SQLi and LFI attacks are driven by the ready availability of free and open-source tools to find vulnerabilities in sites. Such tools include sqlmap (open-source), Acunetix (closed source) and Havij (freeware), as well as a multitude of other tools. The tools are generally meant to be used by legitimate administrators testing their own sites, but are often abused by attackers. Most of these tools can be easily identified by their user agent strings, which most attackers don't bother to change.

These two types of attacks require a very noisy reconnaissance approach. Tools for finding SQL injection vulnerabilities can easily make thousands of requests against a site, testing and probing for an entry point. Blind SQL injection, which amounts to asking a site a series of yes or no questions, can require even more requests.

We have also observed a prevalence of web application scanners. These point-and-shoot tools are easy to obtain and easy to use against any website. They make a high number of requests when looking for SQLi and LFI vulnerabilities. Figure 3-9 shows an example of an LFI attempt as listed in a log.

```
URL: /[redacted]/download.php?f=../../../configuration.php
```

Figure 3-9: An example of what an LFI attempt will look like in a log

The `f` parameter is input to display a specific file on the server in an attempt to force the `configuration.php` file to potentially show server and application configuration settings, which may include usernames and passwords to other systems in the network. To find evidence of LFI attempts in logs, look for the `../` indicator or a changing number of that indicator.

3.6 / WEB APPLICATION SPOTLIGHT: SCRAPERS / A scraper is a specific type of bot whose purpose is to take data from targeted websites, store and analyze it, and make the data available. One example of a scraper is a search engine bot. Other examples are rate aggregators, resellers and SEO analytics services.

For example, Akamai has observed a scraper requesting thousands of articles from news sites. The articles are repurposed for other sites or analyzed to offer SEO services back to the site that was scraped. We have observed other scrapers acquiring store locations, stock prices, and online store inventory. The owners of these scrapers have been traced back to site competitors and companies who advertise services in the business of data analytics.

During Q3, Akamai observed scraping campaigns across multiple industries, but mostly retail. Figure 3-10 shows metrics from a campaign we observed in the retail industry that involved more than 480,000 web scraping attempts to a single site from September 21 through October 20.

Figure 3-10: Akamai observed more than 480,000 web scraping attempts to a single site during a one-month period, September 21 to October 20, 2015

An easy way to identify a scraper is by the request's user agent string. In addition to seeing Googlebot or Bingbot, we see names such as MJ12bot, XoviBot, DotBot, Ahrefsbot, and UptimeRobot. The last one, UptimeRobot, appears to only make HEAD requests to the index page, indicating it is only checking to see if the site is up. Figure 3-11 shows UptimeRobot in action.

```
10/10/2015 07:04...    HEAD    /index.html    Mozilla/5.0+(compatible; UptimeRobot/...
```

Figure 3-11: UptimeRobot makes a HEAD request to check to see if a site is up

Many requests are made to the `robots.txt` and the `index.html` files, and another example we see in the logs is a scraper attempting to find store locations. Figure 3-12 shows a bot using zip codes to brute force a site into giving the desired data.

```
10/10/2015 05:58:0...  GET  failover-[redacted]  /order/storesSearchsearchType=All&zipcode=43716
10/10/2015 05:58:1...  GET  failover-[redacted]  /order/storesSearchsearchType=All&zipcode=43717
10/10/2015 05:58:5...  GET  failover-[redacted]  /order/storesSearchsearchType=All&zipcode=43718
10/10/2015 05:59:1...  GET  failover-[redacted]  /order/storesSearchsearchType=All&zipcode=43719
10/10/2015 05:59:4...  GET  failover-[redacted]  /order/storesSearchsearchType=All&zipcode=43720
```

Figure 3-12: A robot uses zip codes to brute force a site into giving the desired data

An Akamai customer in the automotive sales business received scraping activity using vehicle identification numbers (VINs) to gather data, as shown in Figure 3-13. The site owner received more than 200 million requests in a one-month period.

```
10/10/2015 11:27:24:000  GET  /members/powersearch/vehicleDetails.do  vin=2G1165S31F9107542  Mozilla/5.0 (Windows NT 6.3; WOW64) AppleWebKit/537.
10/10/2015 11:27:24:000  GET  /members/powersearch/vehicleDetails.do  vin=1GCVKREH4FZ251319  Mozilla/5.0 (Windows NT 6.3; WOW64) AppleWebKit/537.
10/10/2015 11:27:24:000  GET  /members/powersearch/vehicleDetails.do  vin=2G1165S39F9104307  Mozilla/5.0 (Windows NT 6.3; WOW64) AppleWebKit/537.
10/10/2015 11:27:24:000  GET  /members/powersearch/vehicleDetails.do  vin=1G11A5SL5FF323982  Mozilla/5.0 (Windows NT 6.3; WOW64) AppleWebKit/537.
10/10/2015 11:27:24:000  GET  /members/powersearch/vehicleDetails.do  vin=1GCRCREC7FZ165246  Mozilla/5.0 (Windows NT 6.3; WOW64) AppleWebKit/537.
10/10/2015 11:27:24:000  GET  /members/powersearch/vehicleDetails.do  vin=1GCVKREH7FZ250374  Mozilla/5.0 (Windows NT 6.3; WOW64) AppleWebKit/537.
10/10/2015 11:27:24:000  POST /members/powersearch/searchResults.do   WT.svl=m_ps_srp_pgnum  Mozilla/5.0 (Windows NT 6.3; WOW64) AppleWebKit/537.
10/10/2015 11:27:25:000  GET  /members/powersearch/vehicleDetails.do  vin=1G11B5SL5FF155760  Mozilla/5.0 (Windows NT 6.3; WOW64) AppleWebKit/537.
10/10/2015 11:27:25:000  GET  /members/powersearch/vehicleDetails.do  vin=5UXKR0C51F0P01546  Mozilla/5.0 (Windows NT 6.3; WOW64) AppleWebKit/537.
10/10/2015 11:27:25:000  GET  /members/powersearch/vehicleDetails.do  vin=1G11B5SL4FF152798  Mozilla/5.0 (Windows NT 6.3; WOW64) AppleWebKit/537.
```

Figure 3-13: Logs show a sample of more than 200 million scraping requests in a one-month period. This example uses a vehicle identification number to gather data

In the requests above, the user agent string is a valid browser. When we look a little further, though, we see the requests are coming from an inexpensive hosting service — this is an indicator that the requests may not be legitimate and are from a scraper.

Our final example is a scraper against the same automotive sales website. The requests shown in Figure 3-14 are from the analytics company Proximic, which is gathering information about automotive inventory.

In many instances we see data analytics companies scraping a website, gathering the information — prices or inventory — and offering the results as an analysis product to companies. Sometimes the analysis is sold back to the companies whose data was scraped.

```
10/12/2015 09:00:37:000  GET  /cars-for-sale/research.xhtml         listingid=41103154&zip=95608&endYear=2010&pageLA...  Mozilla/5.0 (compatible; proximic; +http://www.proximic...
10/12/2015 09:00:40:000  GET  /cars-for-sale/searchresults.xhtml   priceLabel=2000-12000&zip=73119&endYear=2016&pag...  Mozilla/5.0 (compatible; proximic; +http://www.proximic...
10/12/2015 09:00:51:000  GET  /cars-for-sale/searchresults.xhtml   priceLabel=28500+or+less&zip=53221&endYear=2016&p...  Mozilla/5.0 (compatible; proximic; +http://www.proximic...
10/12/2015 09:00:52:000  GET  /cars-for-sale/searchresults.xhtml   &featureCodes=1033,1126,1132&zip=04106&endYear=2...  Mozilla/5.0 (compatible; proximic; +http://www.proximic...
10/12/2015 09:00:53:000  GET  /cars-for-sale/searchresults.xhtml   priceLabel=25000+or+less&zip=98383&endYear=2016&p...  Mozilla/5.0 (compatible; proximic; +http://www.proximic...
10/12/2015 09:00:59:000  GET  /cars-for-sale/searchresults.xhtml   priceLabel=3500+or+less&zip=97862&endYear=2016pa...  Mozilla/5.0 (compatible; proximic; +http://www.proximic...
10/12/2015 09:01:06:000  GET  /cars-for-sale/searchresults.xhtml   priceLabel=50000+or+less&endYear=2016&zip=77833&p...  Mozilla/5.0 (compatible; proximic; +http://www.proximic...
10/12/2015 09:01:21:000  GET  /cars-for-sale/searchresults.xhtml   priceLabel=9000+or+less&zip=94928&endYear=2016&pa...  Mozilla/5.0 (compatible; proximic; +http://www.proximic...
10/12/2015 09:01:34:000  GET  /cars-for-sale/searchresults.xhtml   priceLabel=Any&endYear=2016&pageLayout=list&startY...  Mozilla/5.0 (compatible; proximic; +http://www.proximic...
10/12/2015 09:01:35:000  GET  /cars-for-sale/searchresults.xhtml   priceLabel=9000+or+less&zip=03071&endYear=2010&pa...  Mozilla/5.0 (compatible; proximic; +http://www.proximic...
```

Figure 3-14: Automated requests from the analytics company Proximic gathering information about automotive inventory

3.7 / Methodology / For a long time Akamai has tracked metrics for DDoS attacks; they are typically the most commented on, reprinted, and discussed stats that we produce. Over the years, however, our customers have asked for a similar view into the stealthy web application attacks that plague enterprises, governments and others — the attacks that organizations such as the ***Open Web Application Security Project (OWASP)*** have tracked and ranked according to prevalence and danger.

Figuring out how to give our customers a view of what we see has been a long and arduous challenge. Although Akamai has visibility into 15 – 30% of the world's web traffic, the challenge in meeting this goal has been threefold: how to store the data we see, how to query it, and finally, how to report on it meaningfully.

In the past two years, we've tackled the first two challenges. ***Akamai's Cloud Security Intelligence*** (CSI) platform now stores more than 2 petabytes (PB) of threat intelligence data (2,000 terabytes) — 10 TB of application layer attack data a day for a rolling 30 – 45 days.

Querying the data has taken quite a bit of finesse. To do it, we hired a number of data scientists, analysts and researchers. Today, those researchers make up Akamai's company-wide research team, which has set up dozens of heuristics to automatically query the stored data every hour. The insight they extract from the data feeds improvements to Kona Site Defender and our Client Reputation engine.

The final challenge was reporting on the data. Our reporting methodology is based on some assumptions. First, we divided all Akamai customers into eight verticals. Then, for each of the customers in these eight verticals, we tracked the number of malicious requests across the nine categories of attacks featured in this report during a 12-week period. Next, the frequency of the attack vectors and the accuracy of the signatures used to detect each of the attack categories were given weight in the selection of categories.

As the CSI platform and the capabilities of our team and its resources grow, we look forward to continuing to report on data such as is provided in this report, in addition to exploring new trends as they develop. Please engage us and let us know which types of data you'd like to see in the next report. As long as we can guarantee the anonymity of our customers, we'll continue to share as much as we can in the most useful way possible.

[SECTION]4 AKAMAI EDGE FIREWALL ACTIVITY

This quarter marks the first time Akamai Edge Firewall data is in our security report. Edge Firewall data sets provide a broad look at attack activity at the global platform perimeter — with information on attack traffic coming from 200,000 machines outfitted with Akamai technology.

At the platform perimeter, two dropped packets per second are analyzed, giving us a more accurate, broader look at affected hosts and attack tactics. This data creates a bigger magnifying glass to show what types of non-layer 7 attacks are being attempted against Akamai customers. This report focused on UDP-reflected DDoS attacks, including SSDP, NTP, CHARGEN and Quote of the Day (QOTD).

Among our findings is that the most heavily-abused networks are in China and other Asian countries. While most SSDP attacks tend to be from home connections, NTP, CHARGEN and QOTD are generally from cloud and hosting providers where those services run. We see more repetitive use of the same NTP and CHARGEN reflectors and less reuse of individual SSDP reflectors.

Figure 4-1 is a geographical heat map designed to show the most prevalent areas for the SSDP, CHARGEN, NTP and QOTD attack activity identified in Q3 2015. It is populated by logs identifying more than 1.5 million reflectors. The map shows that the US and Europe are most heavily abused for use as DDoS reflectors.

DDoS Reflector Heatmap, Q3 2015

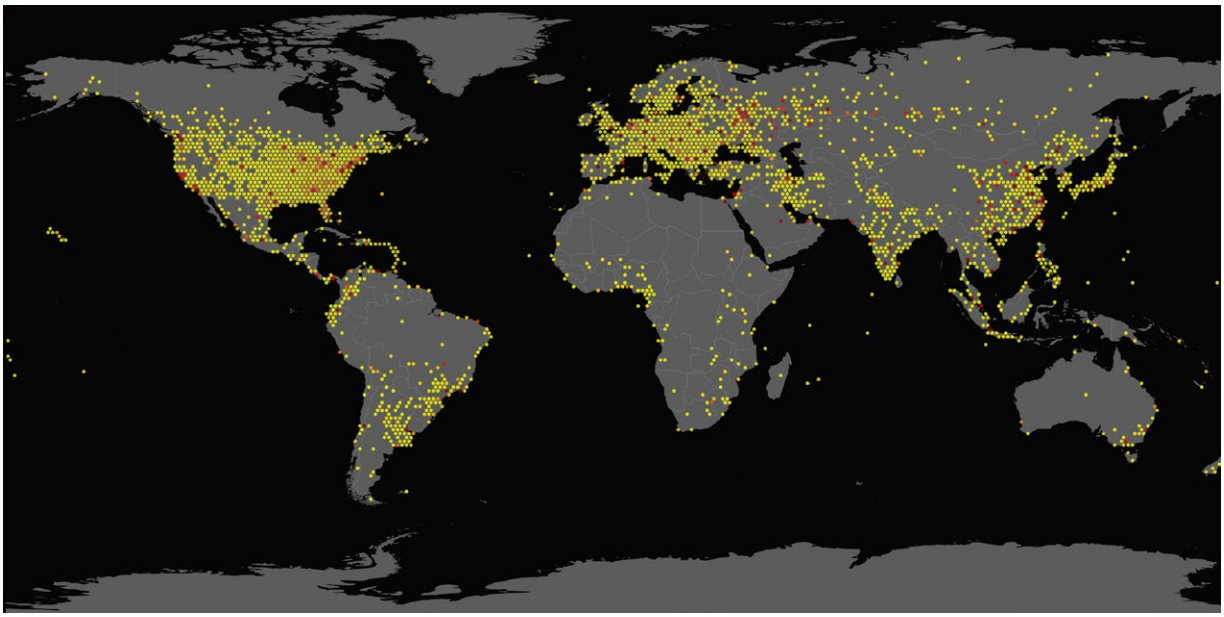

Figure 4-1: A geographical heat map of victim SSDP, CHARGEN, NTP and QOTD reflectors in Q3 2015

As expected, we found SSDP reflectors were used most frequently. This correlates with the earlier data about DDoS attacks and helps explain why SSDP reflection has been a heavily favored vector. Figure 4-2 shows the percentage for each of the four collected reflector types within this platform.

**While NTP accounted for 23% of the reflection sources, only a limited number of these responded in a manner that makes the monlist query a viable amplification source. The number of NTP reflectors that met that criteria was less than the total for CHARGEN.*

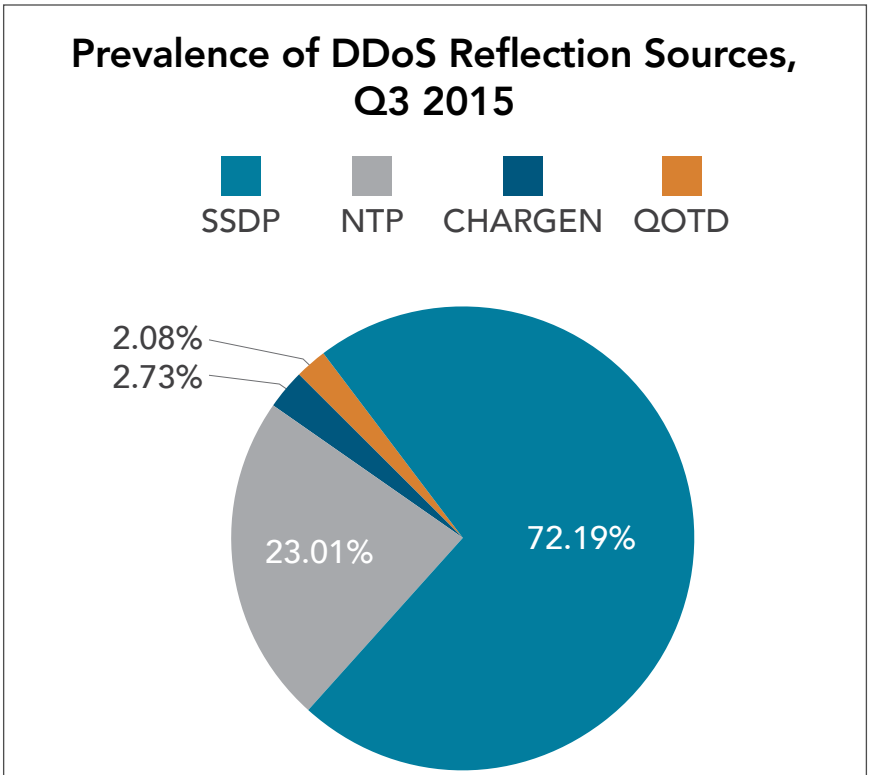

Figure 4-2: The percentage of DDoS reflection sources identified by the Akamai Edge Firewall

The next illustration, Figure 4-3, shows the Top 10 source ASNs for SSDP, CHARGEN, NTP and QOTD reflectors. This further breaks out the data from the geographical heatmap by diving to the origin of the attack traffic. Leading the pack was ASN4837 (CNCGroup China169 Backbone) at 28% of originating traffic, followed by ASN4134 (CHINANET-Backbone) at 25%, and in third place was ASN17676 (Japan Network Information Center) at 15%. The other ASN and associated Internet providers sourced 3 – 6%.

Top 10 DDoS Reflection Source ASNs, Q3 2015

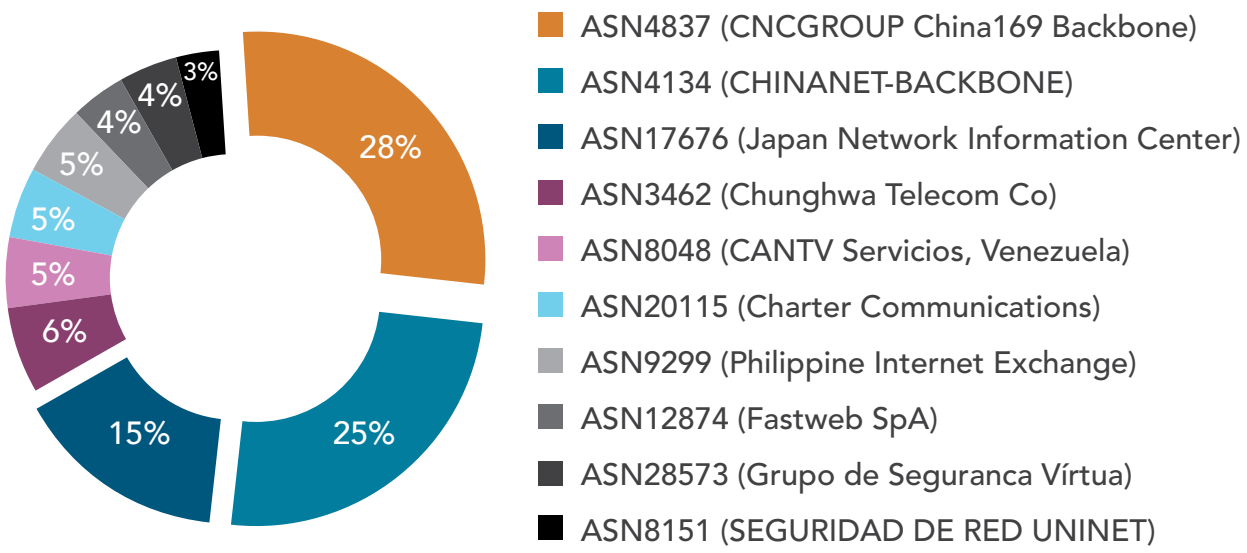

Figure 4-3: The top origins of attacks that use DDoS reflection methods are two Chinese networks and one Japanese network

[SECTION]5 CLOUD SECURITY RESOURCES

Akamai released five threat advisories and attack case studies in Q3 2015, as summarized here.

5.1 / *New DDoS Reflection Techniques* / For malicious actors looking to bring a website or service offline, distributed reflection denial of service attacks have been a popular weapon of choice. Going into Q3 2015, the reflection-based attack trend continued.

```
16:01:14.572122 IP (tos 0x0, ttl 123, id 24986, offset 0, flags [none], proto UDP (17),
length 1460)
x.x.x.x.5093 > x.x.x.x.1351: UDP, length 1432
16:01:14.576212 IP (tos 0x0, ttl 123, id 24988, offset 0, flags [none], proto UDP (17),
length 1460)
x.x.x.x.5093 > x.x.x.x.1351: UDP, length 1432
```

Figure 5-1: Sentinel reflection attack signature

In a reflection DDoS attack, a malicious actor begins by sending a query to a victim IP address. The victim is an unwitting accomplice in the attack. The victim could be any device on the Internet that exposes a reflectable UDP service. The attacker's query is spoofed to appear to originate from the attacker's target. The attacker uses an automated attack tool to send malicious queries at high rates to a large list of victims, who will in turn respond to the target.

In a ***case study released in October***, Akamai outlined three new reflected DDoS attacks: NetBIOS name server (NBNS) reflection, RPC portmap reflection, and Sentinel reflection, which reflects off of licensing servers. Akamai has mitigated each of these reflection attack methods multiple times while protecting our customers from DDoS attacks.

One of the 10 reflection attack campaigns was especially large. The RPC reflection attack vector was used in a mega DDoS attack that generated more than 100 Gbps.

NetBIOS reflection DDoS attacks were observed by Akamai as occurring sporadically from March to July 2015. Although legitimate and malicious NBNS queries to UDP port 137 are a common occurrence, a response flood was first detected in March 2015 during a mitigated DDoS attack.

5.2 / *XOR DDoS* / In September, Akamai's Security Intelligence Response Team (***SIRT***) tracked XOR DDoS, a Trojan malware attackers were using to hijack Linux machines for participation in a DDoS botnet. The bandwidth of DDoS attacks

coming from the XOR DDoS botnet had ranged from a few Gbps to 150+ Gbps. The gaming sector had been the primary target, followed by educational institutions. Akamai SIRT released an ***XOR DDoS threat advisory*** in late September.

At the time, the botnet was attacking up to 20 targets per day, 90% of which were in Asia. Akamai mitigated two DDoS attacks orchestrated by the XOR DDoS botnet on the weekend of August 22. The largest attack of the quarter — which peaked at 149 Gbps — was launched by the XOR DDoS botnet. Additional details about that attack are shown in Figure 5-2.

Akamai Scrubbing Center	Peak Gbps	Peak Mpps
Hong Kong	25.66	5.50
Washington	19.70	3.40
San Jose	18.45	3.80
Frankfurt	24.00	3.50
London	11.00	4.30
Tokyo	50.50	15.70

Figure 5-2: Traffic distribution by Akamai DDoS scrubbing center for a SYN flood orchestrated by the XOR botnet

Late in the quarter, the same botnet was the source of a DNS flood against a customer using Akamai's FastDNS infrastructure. The attack campaign started with a DNS flood of 30 Mpps and escalated into a SYN flood ramping up to 140 Gbps and more than 75 Mpps.

XOR DDoS is an example of attackers building botnets from Linux systems instead of Windows-based machines. A decade ago, Linux was perceived as a more secure alternative to Windows systems, which suffered the most attacks at the time. As a result, companies increasingly adopted Linux as part of their security-hardening efforts. But Linux offers no guarantees. The malware does not spread via a host vulnerability. Rather, it populates via Secure Shell (SSH) services that are susceptible to brute-force attacks due to weak passwords.

Other recent examples of Linux-based malware include the ***Spike DDoS toolkit*** (which also targeted Windows machines) and ***IptabLes and IptabLex malware***. There are an increasing number of Linux vulnerabilities for malicious actors to target, such as the ***heap-based buffer overflow vulnerability*** found earlier this year in the GNU C library. However, XOR DDoS itself does not exploit a specific vulnerability.

XOR DDoS captured the attention of technology news outlets, including SC Magazine, which ***describes attacks*** that alter installations based on the victim's Linux environment. A rootkit is also deployed to cloak the main attack. ***The Avast blog*** has also focused on XOR DDoS attacks.

Akamai SIRT's research indicated the malware was of Asian origin, based on the command-and-control (C2) IP addresses and source IP addresses of the attack payloads.

We expect XOR DDoS activity to continue as attackers refine their methods. Further, we anticipate a more diverse selection of DDoS attack types in future versions of the malware. XOR DDoS malware is part of a wider trend of which companies must be aware: Attackers are targeting poorly configured and unmaintained Linux systems for use in botnets and DDoS campaigns.

5.3 / *More Attack Activity from DD4BC* and the Rise of Armada Collective / In early September, Akamai ***released a case study*** about the DDoS attacks from the bitcoin extortionist group DD4BC, based on Akamai SIRT's observation of attack traffic directed at our customers from September 2014 through August 2015. This is shown in Figure 5-3. The team identified 114 DD4BC attacks, including more aggressive measures that target brand reputation through social media.

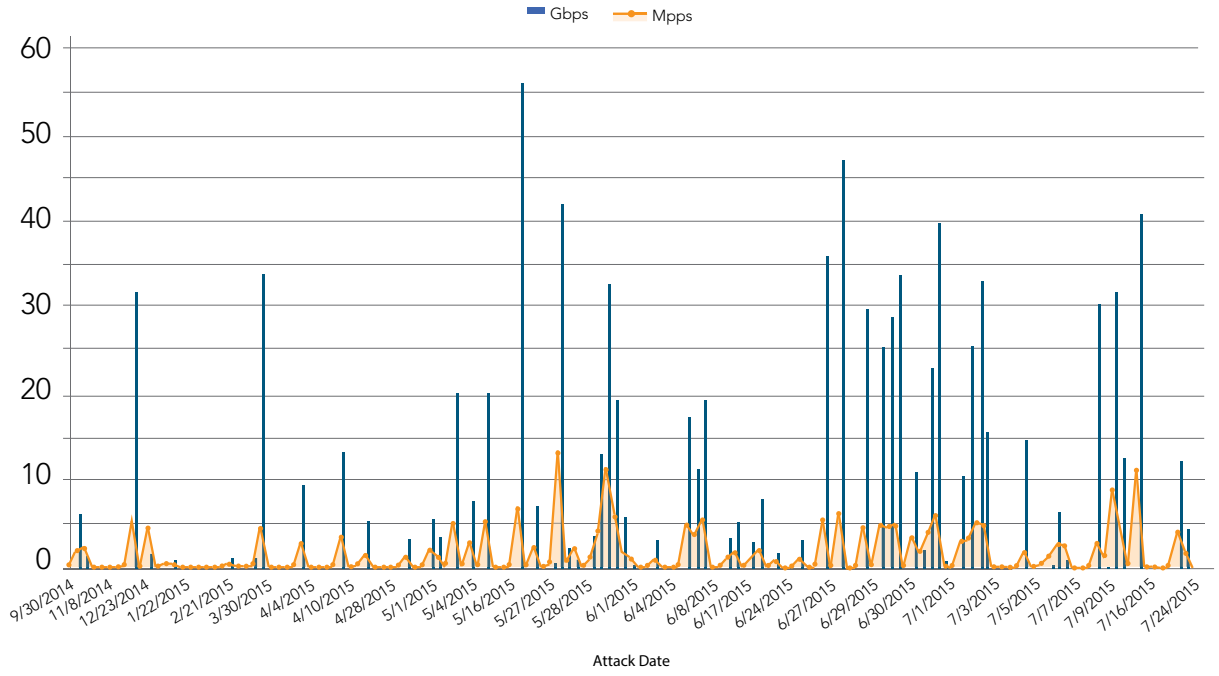

Figure 5-3: Attack timeline of bandwidth and packets per second for DD4BC events

DD4BC has been using the threat of DDoS attacks to secure bitcoin payments from its victims for protection against future attacks. The latest attacks — focused primarily on the financial service industry — involved new strategies and tactics intended to harass, extort and ultimately embarrass the victim publicly.

By the end of Q3 2015, Akamai began seeing attacks from a group calling itself the Armada Collective. The initial thinking was that DD4BC had changed its name, but on further inspection it appeared more likely that Armada Collective was a copycat group. Like DD4BC, the Armada Collective threatens a victim with emails claiming that a DDoS attack is forthcoming against their victim's unless they pay a specified amount in bitcoins. In one case, the group demanded that an Akamai customer pay 30 bitcoins or all their servers would be DDoSed.

Armada Collective claimed that their DDoS attacks were very powerful — that they had the power to launch an attack of more than 1 Terabit per second (Tbps). At the time of writing, the largest Armada Collective attack mitigated by Akamai peaked at 771 Mbps, less than 1 Gbps.

5.4 / *Cloudpiercer Discovery Tool* / Late in Q3, researchers ***released details*** of a tool that allowed users to discover origin servers. The tool was named Cloudpiercer, and it used a number of techniques to locate the IP addresses of origin servers.

Cloudpiercer bundled several previously known methods with some stated new ones to simplify reconnaissance against targets. It is a reconnaissance tool, not an attack tool. A potential attacker may use similar methods to search for a customer's datacenter IP addresses or netblock(s), but will have to use other services or technologies to perform an actual DDoS or web application attack.

Cloudpiercer required verification of ownership of a site before it could be tested, which limited the ability of attackers to use the tool. However, the methods of discovery described in the paper might allow an adversary to recreate a tool without verification steps.

Akamai's SIRT analyzed the methods used by the tool and found no proof of a large number of documented discovery cases using these types of techniques. The security community had been aware of these methods for several years.

5.5 / *CDN Vulnerability Unveiled at Black Hat USA 2015* / On August 6 at the Black Hat Security Conference in Las Vegas, ***Bishop Fox***, a security research and penetration testing firm, announced the discovery of a vulnerability that allows an outside actor to conduct a cross-site request forgery (CSRF)/Server-Side Request Forgery (SSRF) attack using a combination of exploits.

This vulnerability relied on the Akamai platform in two ways: specially crafted legacy resource locators (also called v1 ARLs) in combination with specific versions of ***Flow Player***.

Ahead of the announcement, Akamai closed the vulnerability by disabling the use of v1 ARLs to go forward to mediapm.edgesuite.net. In addition, Akamai made changes to protect customers using the related Multi-Domain Config feature and continues to make security improvements surrounding other uses of v1 ARLs on our platform.

The researchers who discovered this vulnerability coordinated closely with Akamai to identify exposed domains prior to public release. Thanks to their cooperation, Akamai was already communicating with customers it believed had been exposed to the vulnerability, informing them of remediation plans. To date, there is no evidence that indicates this CSRF was used maliciously.

5.6 / ***Another OpenSSL Vulnerability*** **/** In July, Akamai was made aware of an OpenSSL vulnerability addressed in ***OpenSSL versions 1.0.2d and 1.0.1p***. Akamai does not use the vulnerable versions of OpenSSL and was therefore not affected.

The vulnerability was reported to OpenSSL on June 24. The fix was developed by the BoringSSL project, and released by OpenSSL on July 9.

Though Akamai is unaffected, we recommended that sites running OpenSSL in their origin infrastructure consult their security advisory team to review the vulnerability, upgrade software and address the vulnerability as necessary.

[SECTION]6
LOOKING FORWARD

In the coming months, we expect more records to be set for the number of DDoS attacks recorded on Akamai's routed network, though the attack vectors and methods will continue to vary.

Now that we're able to provide analysis of traffic based on the ASN assigned to traffic in association with its BGP routing, readers can expect a sharpening focus. We expect the US to remain a top source of malicious traffic because of the sheer

number of devices, vulnerabilities and users in the US, and it is likely that cloud providers will remain the biggest trouble spot unless they do more to improve their internal security procedures.

Though activity from DD4BC appears to have quieted, attacks from copycats like the Armada Collective will probably continue.

Distributed reflection denial of service attacks will remain a popular weapon of choice for attackers, though it remains to be seen if NetBIOS, RPC portmap and Sentinel will remain popular reflection DDoS attacks. Surprisingly, despite a decreasing number of available resources, NTP reflection surged near the end of Q3 2015 and continues into Q4.

Expect the heavy barrage of DDoS attacks against the gaming industry to continue, as players keep looking for an edge over competitors, and security vulnerabilities in gaming platforms continue to attract attackers looking for low-hanging fruit. Financial services will also remain a top target given the myriad opportunities malicious actors have to extract and monetize sensitive data.

We will also continue to see malware in ads and third-party service attacks as attackers continue to find security holes in the many widgets and plugins used across myriad platforms.

We expect retailers to continue to suffer the vast majority of web application attacks given the potential financial gains for attackers, and that SQLi and LFI will remain favorite vectors, because free and open-source tools are plentiful to find these vulnerabilities in sites.

Collaboration continues to be an imperative for the software and hardware development industry, application and platform service providers, and the security industry in order to break the cycle of mass exploitation, botnet building and monetization.

Data Sources / The data in this report is based on attacks observed and identified by Akamai. The trends are affected in various ways, not all of which are directly related to increases in attack activity. Example factors include changes in the distribution of our customer base, the launch of new products, and improvements to attack sensors.

The Akamai platform consists of more than 200,000 servers in more than 100 countries around the globe and regularly transmits between 15 – 30% of all Internet traffic. In February 2014, Akamai added the Prolexic routed network to its portfolio, a resource specifically designed to fight DDoS attacks. This report draws its data from the two platforms in order to provide information about current attacks and traffic patterns around the globe.

The Akamai Intelligent Platform protects customers by being massively distributed, protected by the use of the Kona Site Defender and the ability to absorb attack traffic closest to where it originates. In contrast, the routed DDoS solution protects customers by routing traffic to scrubbing centers where experienced incident responders use a variety of tools to remove malicious traffic before passing clean traffic to origin servers. The two types of technology are complementary and provide two lenses through which we can examine traffic on the Internet.

Content
David Fernandez, Editor in Chief
Bill Brenner, Managing Editor
Jose Arteaga, Data Visualization and Research
Ezra Caltum, Web Application Threat Research
Martin McKeay, Senior Editor

Contributors
Dave Lewis, Editor
Jon Thompson, Threat Data Modeling
Patrick Laverty, Research
Larry Cashdollar, Research
Chad Seaman, Research

Design
Shawn Doughty, Creative Direction
Brendan O'Hara, Art Direction/Design

Contact
stateoftheinternet@akamai.com
Twitter: *@State_Internet* / *@akamai*
www.stateoftheinternet.com

About Akamai* As the global leader in Content Delivery Network (CDN) services, Akamai makes the Internet fast, reliable and secure for its customers. The company's advanced web performance, mobile performance, cloud security and media delivery solutions are revolutionizing how businesses optimize consumer, enterprise and entertainment experiences for any device, anywhere. To learn how Akamai solutions and its team of Internet experts are helping businesses move faster forward, please visit *www.akamai.com* or *blogs.akamai.com*, and follow *@Akamai* on Twitter.

Akamai is headquartered in Cambridge, Massachusetts in the United States with operations in more than 57 offices around the world. Our services and renowned customer care are designed to enable businesses to provide an unparalleled Internet experience for their customers worldwide. Addresses, phone numbers and contact information for all locations are listed on *www.akamai.com/locations*.
