[Volume 3 / Number 1]

akamai's [state of the internet] / security

Q1 2016 report

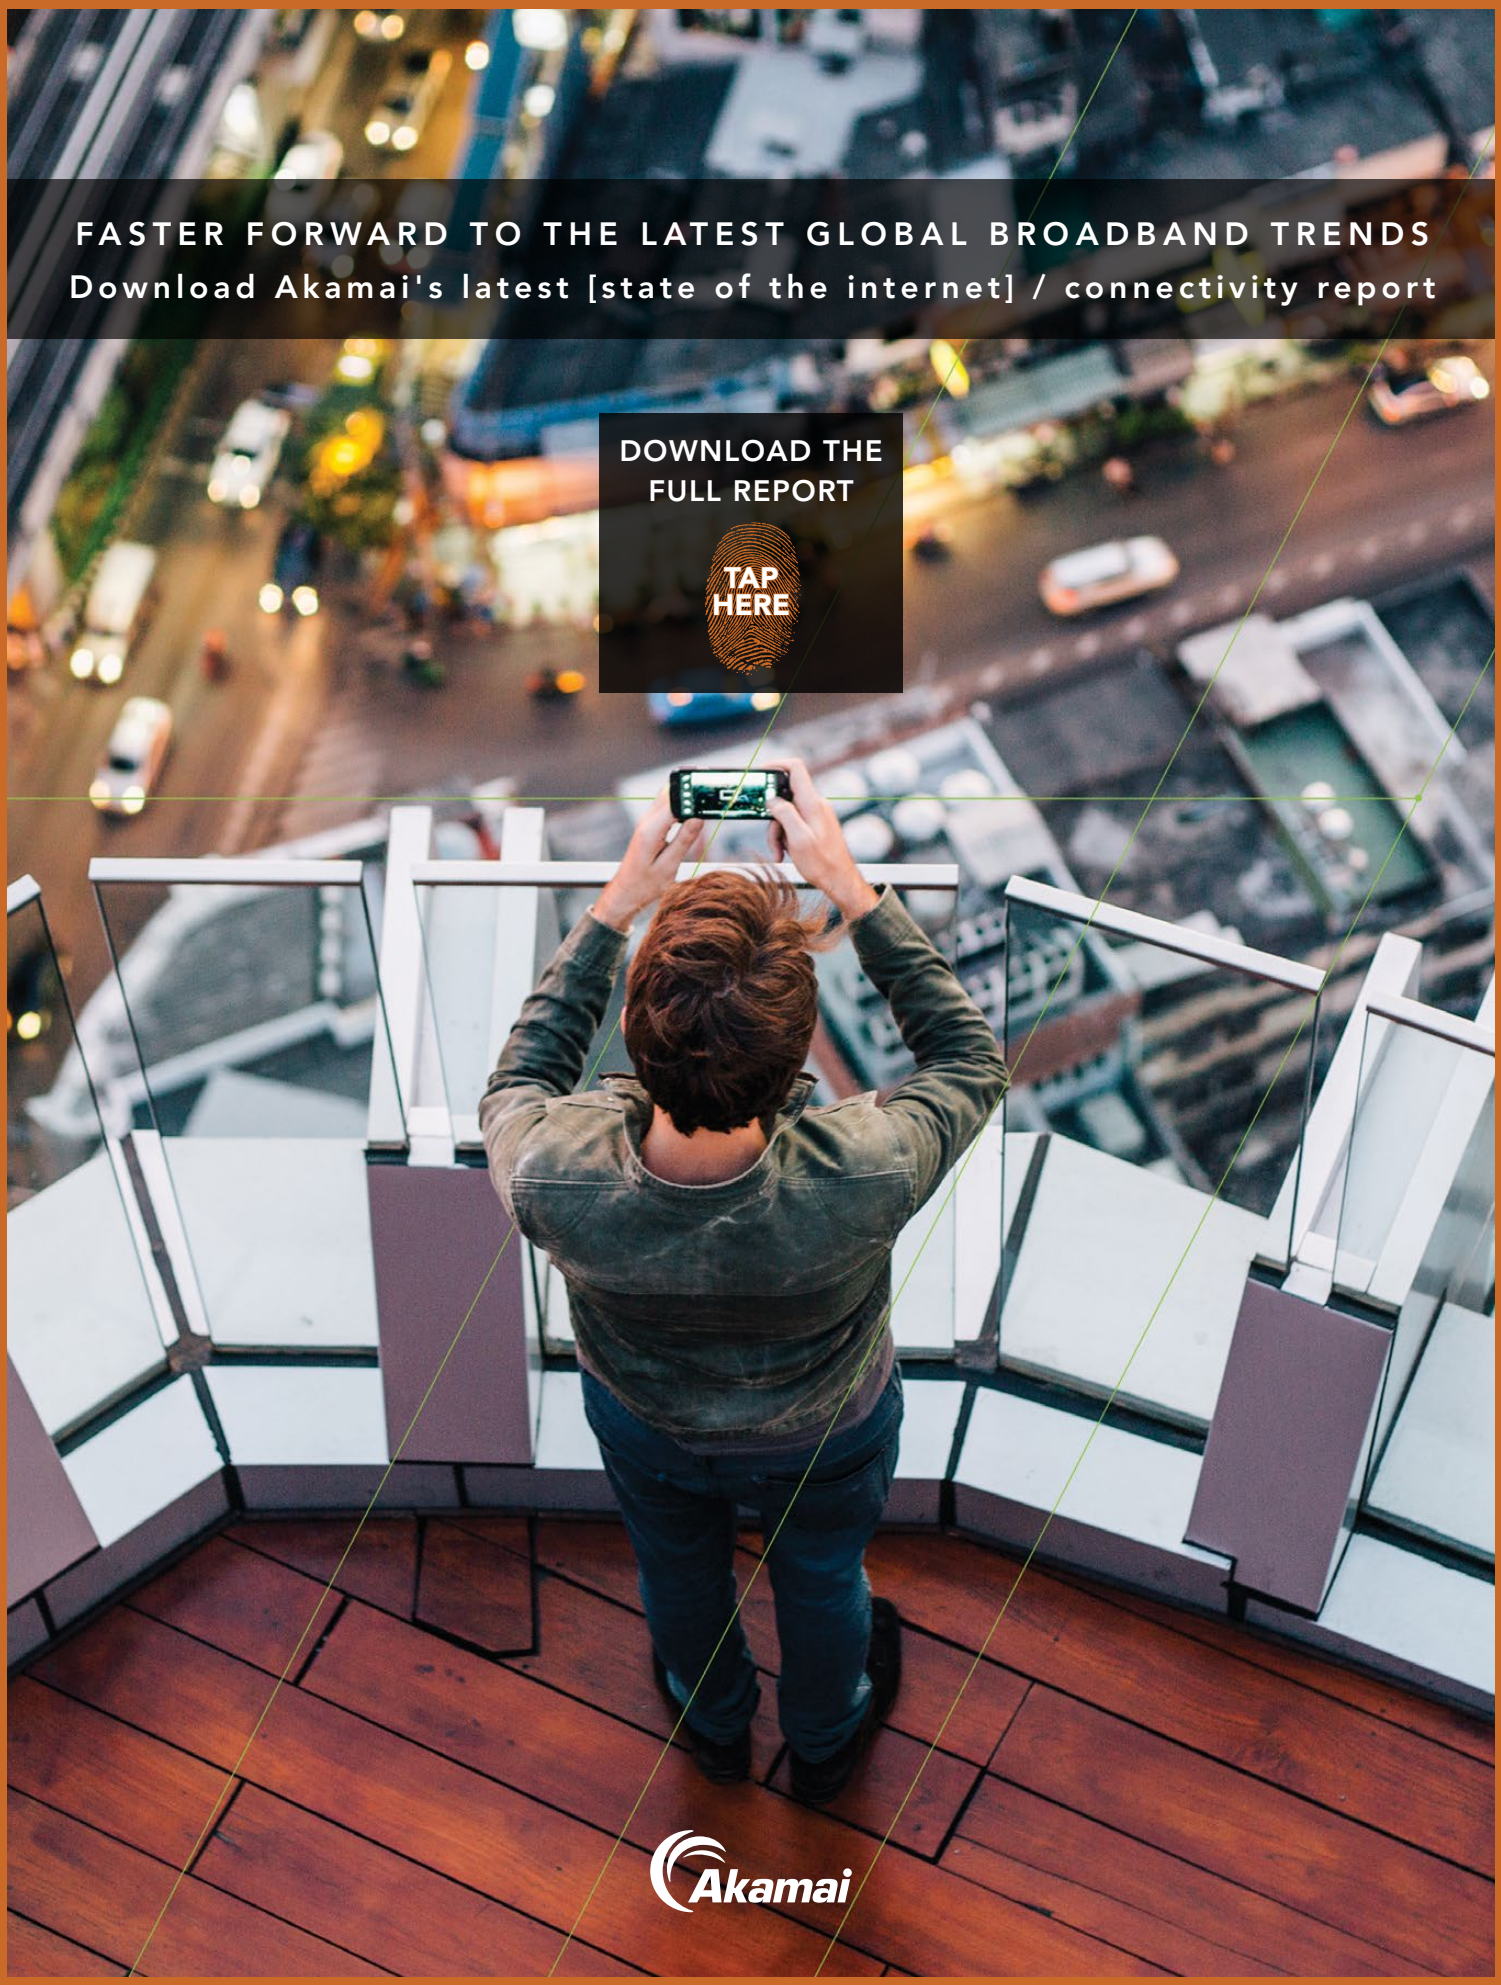
FASTER FORWARD TO THE LATEST GLOBAL BROADBAND TRENDS
Download Akamai's latest [state of the internet] / connectivity report
DOWNLOAD THE FULL REPORT
TAP HERE
Akamai

LETTER FROM THE EDITOR / The *Q1 2016 State of the Internet / Security Report* combines data from Akamai's global infrastructure and the routed DDoS solution.

The Akamai Intelligent Platform™ protects customers by being massively distributed, providing several cloud security solutions, and having the ability to absorb attack traffic closest to its origin. Akamai's Cloud Security Intelligence (CSI) data analysis engine stores more than 2 petabytes (PB) of threat intelligence data. This includes 10 terabytes (TB) of application-layer attack data per day, for a rolling 30 – 45 days. We have more than 70 heuristics to automatically query the stored data every hour. The insight we extract from the data feeds improvements to cloud security solutions and our client intelligence engine.

The routed DDoS solution is designed to protect customers by routing traffic to our global scrubbing centers where experienced incident responders use a variety of mitigation and monitoring tools to remove malicious traffic before passing clean traffic to the customer network.

Each solution collects a distinct data set that represents a unique view of the Internet, allowing us to compare different indicators of attack activity.

The data in this report is based on attacks observed and mitigated by Akamai. The attack trends are affected by several factors, including increases in attack activity, changes in the distribution of our customer base, the launch of new products, and improvements to attack sensors.

Our threat landscape overview comprises an extensive review of the data, allowing us to explore which industries among our customer base suffered the highest attack volumes, which attack techniques and vectors were most common, where malicious traffic originated, and how attack trends evolved.

In this quarter's issue, we have included new DDoS and web application attack visualizations, along with a new dataset examining bot and scraper interactions with Akamai's Intelligent Platform.

The report authors include security professionals from several divisions within Akamai, including the Akamai Security Intelligence Response Team (SIRT), the Threat Research Unit, InfoSec, and the Custom Analytics group. We hope you find the report valuable.

Thank you.

— Akamai's *State of the Internet / Security* Team

If you have comments, questions, or suggestions regarding the *State of the Internet / Security Report*, connect with us via email at ***SOTISecurity@akamai.com***. You can also interact with us in the *State of the Internet / Security* subspace on the Akamai Community at ***https://community.akamai.com***. For additional security research publications, please visit us at ***www.akamai.com/cloud-security***.

6 AT A GLANCE

7 [SECTION][1] = EMERGING TRENDS

11 [SECTION][2] = DDoS ACTIVITY
12 2.1 / DDoS Attack Vectors
18 2.2 / Mega Attacks
20 2.3 / DDoS Attack Source Countries
23 2.4 / DDoS Attacks by Industry
25 2.5 / DDoS Attacks— A Two-Year Look Back
28 2.6 / Reflection DDoS Attacks, Q1 2015–Q1 2016
30 2.7 / Repeat DDoS Attacks by Target
31 2.8 / DDoS Attack Spotlight

35 [SECTION][3] = WEB APPLICATION ATTACK ACTIVITY
36 3.1 / Web Application Attack Vectors
39 3.2 / Top 10 Source and Target Countries
42 3.3 / Web Application Attacks by Industry
45 3.4 / A Normalized View of Web Application Attacks by Industry
52 3.5 / Bot Traffic Analysis
55 3.6 / Web Application Security Spotlight: Attack Takeover Campaign

60 [SECTION][4] = AKAMAI INTELLIGENT PLATFORM™ FIREWALL ACTIVITY

65 [SECTION][5] = CLOUD SECURITY RESOURCES
65 5.1 / Scraper and Bot Series— When Good Bots Go Bad
67 5.2 / #OpKillingBay Expands Attacks
68 5.3 / BillGates Malware Used in DDoS Attacks
68 5.4 / Akamai Responds to Forwarding-Loop Issue
69 5.5 / IKE/IKEv2 Ripe for DDoS Abuse
70 5.6 / Akamai and the Glibc Vulnerabilitty (CVE-2015-7547)
70 5.7 / Akamai and the DROWN Vulnverability
71 5.8 / DNSSEC Targeted in DNS Reflection, Amplification DDoS Attacks
72 5.9 / Akamai Customers Not Vulnerable to SLOTH
73 5.10 / How Web Applications Become SEO Pawns

74 [SECTION][6] = LOOKING FORWARD

AT A GLANCE

DDoS attacks, Q1 2016 vs. Q1 2015
125.36% increase in total DDoS attacks
142.14% increase in infrastructure layer (layers 3 & 4) attacks
34.98% decrease in the average attack duration: 16.14 vs. 24.82 hours
137.5% increase in attacks > 100 Gbps: 19 vs. 8

DDoS attacks, Q1 2016 vs. Q4 2015
22.47% increase in total DDoS attacks
23.17% increase in infrastructure layer (layers 3 & 4) attacks
7.96% increase in the average attack duration: 16.14 vs. 14.95 hours
280% increase in attacks > 100 Gbps: 19 vs. 5

Web application attacks, Q1 2016 vs. Q4 2015
25.52% increase in total web application attacks
1.77% decrease in web application attacks over HTTP
235.99% increase in web application attacks over HTTPS
87.32% increase in SQLi attacks

What you need to know

- Akamai mitigated 4,523 DDoS attacks within Akamai's routed DDoS mitigation platform—an increase from the previous quarter's 3,693 attack events.
- A record 19 attacks exceeded 100 Gbps, with the largest attacks hitting the software & technology, gaming, and media & entertainment sectors. That's a significant increase from the five DDoS attacks that registered more than 100 Gbps in Q4.
- The continued rise of multi-vector DDoS attacks in Q1 suggests that attack tools have continued to grow more sophisticated—causing problems for security practitioners, since each attack vector requires unique mitigation controls.
- Many copycat entities are imitating the tactics of attack groups like Armada Collective and DD4BC, sending victims ransom messages promising large-scale DDoS unless they pay a specified amount of bitcoins.
- NTP reflectors accounted for 59% of reflection DDoS attacks. The use of NTP reflectors increased 72% compared with Q4 2015.
- Web application attacks increased nearly 26% compared with Q4 2015. As in past quarters, the retail sector remained the most popular attack target, receiving 43.4% of the attacks.
- In a shift from last quarter, we saw a 2% decrease in web application attacks over HTTP and a 236% increase in web application attacks over HTTPS. There was also an 87% increase in SQLi attacks compared with the previous quarter.
- In Q1 2016, the US was the main source of web application attacks, accounting for 43% of attack origin traffic.
- One customer was targeted by 4% or more of the total web application attacks in seven out of the nine vectors, while another customer suffered 12% of the total number of attacks observed in the whole quarter and two customers together accrued 60% of the Shellshock attacks.

[SECTION][1] EMERGING TRENDS

In Q1 2016, Akamai witnessed 4,523 attack events across our routed solution, one of three networks used to protect customers against Distributed Denial of Service (DDoS) attacks. That represents a significant increase from the previous quarter's 3,693 attack events. This increase was largely driven by repeat attacks on customers rather than a broadening of the number of targets.

In all, 59% of the DDoS attacks mitigated in Q1 2016 were multi-vector attacks, building on a trend we saw last quarter.

We continued to see a rise in the use of stresser/booter-based botnets by attackers. These tools rely upon the use of reflection attacks, bouncing traffic off servers running vulnerable services such as DNS, CHARGEN, and NTP. Because this type of attack depends on large packet sizes to increase attack bandwidth, it consequently reduces the average number of packets per attack. In other words, while the average gigabits per second (Gbps) per attack increased, the average number of packets per second (PPS) decreased.

In Q1 2016, a record 19 attacks exceeded 100 Gbps, the largest hitting the software & technology, gaming, and media & entertainment sectors. That's a significant increase from the five DDoS attacks that registered more than 100 Gbps in Q4. In comparison, we saw eight in Q3 2015 and 17 in Q3 2014 — the previous record.

The largest DDoS attack measured 289 Gbps, a 20-point drop over the largest attack in the previous quarter (309 Gbps). More interesting is that, aside from the largest attack, most of the mega-attacks seemed to use tools common to booters/stressers.

Six DDoS attacks in Q1 exceeded 30 million packets per second (Mpps) and two attacks peaked at more than 50 Mpps. The packet rate can affect some routers and networks more than the number of bytes per packet because even the smallest packets consume memory, tying up resources. This resource consumption can result in packet loss within these routers and potentially cause collateral damage.

In terms of DDoS attack sources, China took the lead with 27%, the US came in second with 17%, and Turkey was third with 10% of attacking IP addresses. This was the second consecutive quarter that Turkey made the top 10 list.

The online gaming sector was hit hard with DDoS attacks in Q1 2016, accounting for 55% of all DDoS attacks. This was no surprise, since gaming has been the hardest-hit by DDoS attacks for the last several years. The frequency of attacks in that industry was about the same as in Q4 2015.

Of the 24 DDoS attack vectors tracked this quarter, four — UDP fragment, NTP, DNS, and CHARGEN — made up nearly 70% of the attacks.

Twenty-seven percent of DDoS attacks contained UDP fragments in Q1 2016, a six-point jump over the previous quarter. Some of this was a direct result of the amplification factor included in reflection-based attacks, primarily from the CHARGEN, DNS, and SNMP protocols, all of which have potentially large payloads.

Web application attacks increased nearly 26% compared with Q4 2015. As in past quarters, the retail sector remained the most popular attack target, receiving 43% of the attacks. Hotel & travel suffered 13% of attacks while financial services suffered 12% of attacks, followed by high technology (9%), media & entertainment (7%), the public sector (3%), Software-as-a-Service (SaaS) (3%), and business services (2%).

In a shift from last quarter, we saw a 2% decrease in web application attacks over HTTP and a 236% increase in web application attacks over HTTPS. There was also an 87% increase in SQLi attacks compared with the previous quarter.

In Q1 2016, the US was the main source of web application attacks, accounting for 43% of attack origin traffic. Brazil was the second-largest source country at 12%, followed by China and the Netherlands (8% each), Romania (7%), Russia and the UK (6% each), Germany (4%), and the Ukraine and India (3% each).

When the attack source was the US, the main attack targets were in the retail industry, followed by manufacturing and media. In those cases, the preferred attack methods were SQLi, LFI, and RFI. A big difference with attack sources from Brazil was that the main destinations were not only the US, but also India and Australia.

In recent months, a global and respectable cloud Infrastructure-as- a-Service (IaaS) provider opened data centers in Brazil. Since the opening of these data centers, Akamai has seen a large increase in the amount of malicious traffic coming out of Brazil — specifically from the aforementioned data centers. Most of those attacks were against a Brazilian customer in the retail industry.

When the attacks originated in Russia, the destinations were mostly in the retail industry in the US and the UK.

Akamai published 10 threat advisories, vulnerability updates and attack case studies in Q1. They include:

- ***Scraper and Bot Series — When Good Bots Go Bad***
- ***#OpKillingBay Expands Attacks***
- ***BillGates Malware Used in DDoS Attacks***
- ***Akamai Responds to Forwarding-Loop Issue***
- ***IKE/IKEv2 Ripe for DDoS Abuse***
- ***Akamai and the Glibc Vulnerability (CVE-2015-7547)***
- ***Akamai and the DROWN Vulnerability***
- ***DNSSEC Targeted in DNS Reflection, Amplification DDoS Attacks***
- ***Akamai Customers Not Vulnerable to SLOTH***
- ***How Web Applications Become SEO Pawns***

[SECTION]2
DDoS ACTIVITY

Compared with the same period a year ago, Q1 2016 saw a 126% increase in total DDoS attacks and a 142% increase in infrastructure layer (layers 3 & 4) attacks. The average duration of attacks this quarter was 16.14 hours, a 35% decrease from the 24.82 hours we saw in Q1 2015. Average peak bandwidth increased 138% over the same period last year, with 19 attacks pushing 100 Gbps compared with eight in Q1 2015.

The quarter's activity reflected the ongoing trend where attackers make use of cheap, easily-accessed booter/stresser tools, which use reflection attack techniques instead of directly generating their own payloads. The high number of mega attacks this quarter also seemed to originate from these tools. In the past, very few attacks generated with booter/stresser tools exceeded the 100 Gbps mark.

2.1 / DDoS Attack Vectors / As shown in Figure 2-1, of the 24 DDoS attack vectors tracked this quarter, four — UDP Fragment, NTP, DNS, and CHARGEN — made up nearly 70% of the attacks.

Twenty-seven percent of DDoS attacks contained UDP fragments in Q1 2016, a six-point jump over the previous quarter.

Some of this was a direct result of the amplification factor included in reflection-based attacks, primarily from the CHARGEN, DNS, and SNMP protocols, all of which have potentially large payloads.

It was also common to see attack campaigns where UDP flood payloads were set to exceed the default Maximum Transmission Unit (MTU) size of 1,500 bytes, which results in a greater likelihood of fragmentation and overhead for the target. Payload size is configurable in many DDoS attack tools.

Occasionally, we even saw attacks where the packet size was set to 65,000+ bytes. Of course, for these attacks, the attack source is also consuming large amounts of outbound bandwidth, so it is limited by the bandwidth available at the attacking machine's ISP or hosting provider.

Thirty percent of attacks used NTP or DNS, while SYN floods represented 7% of attacks, compared with 10% last quarter.

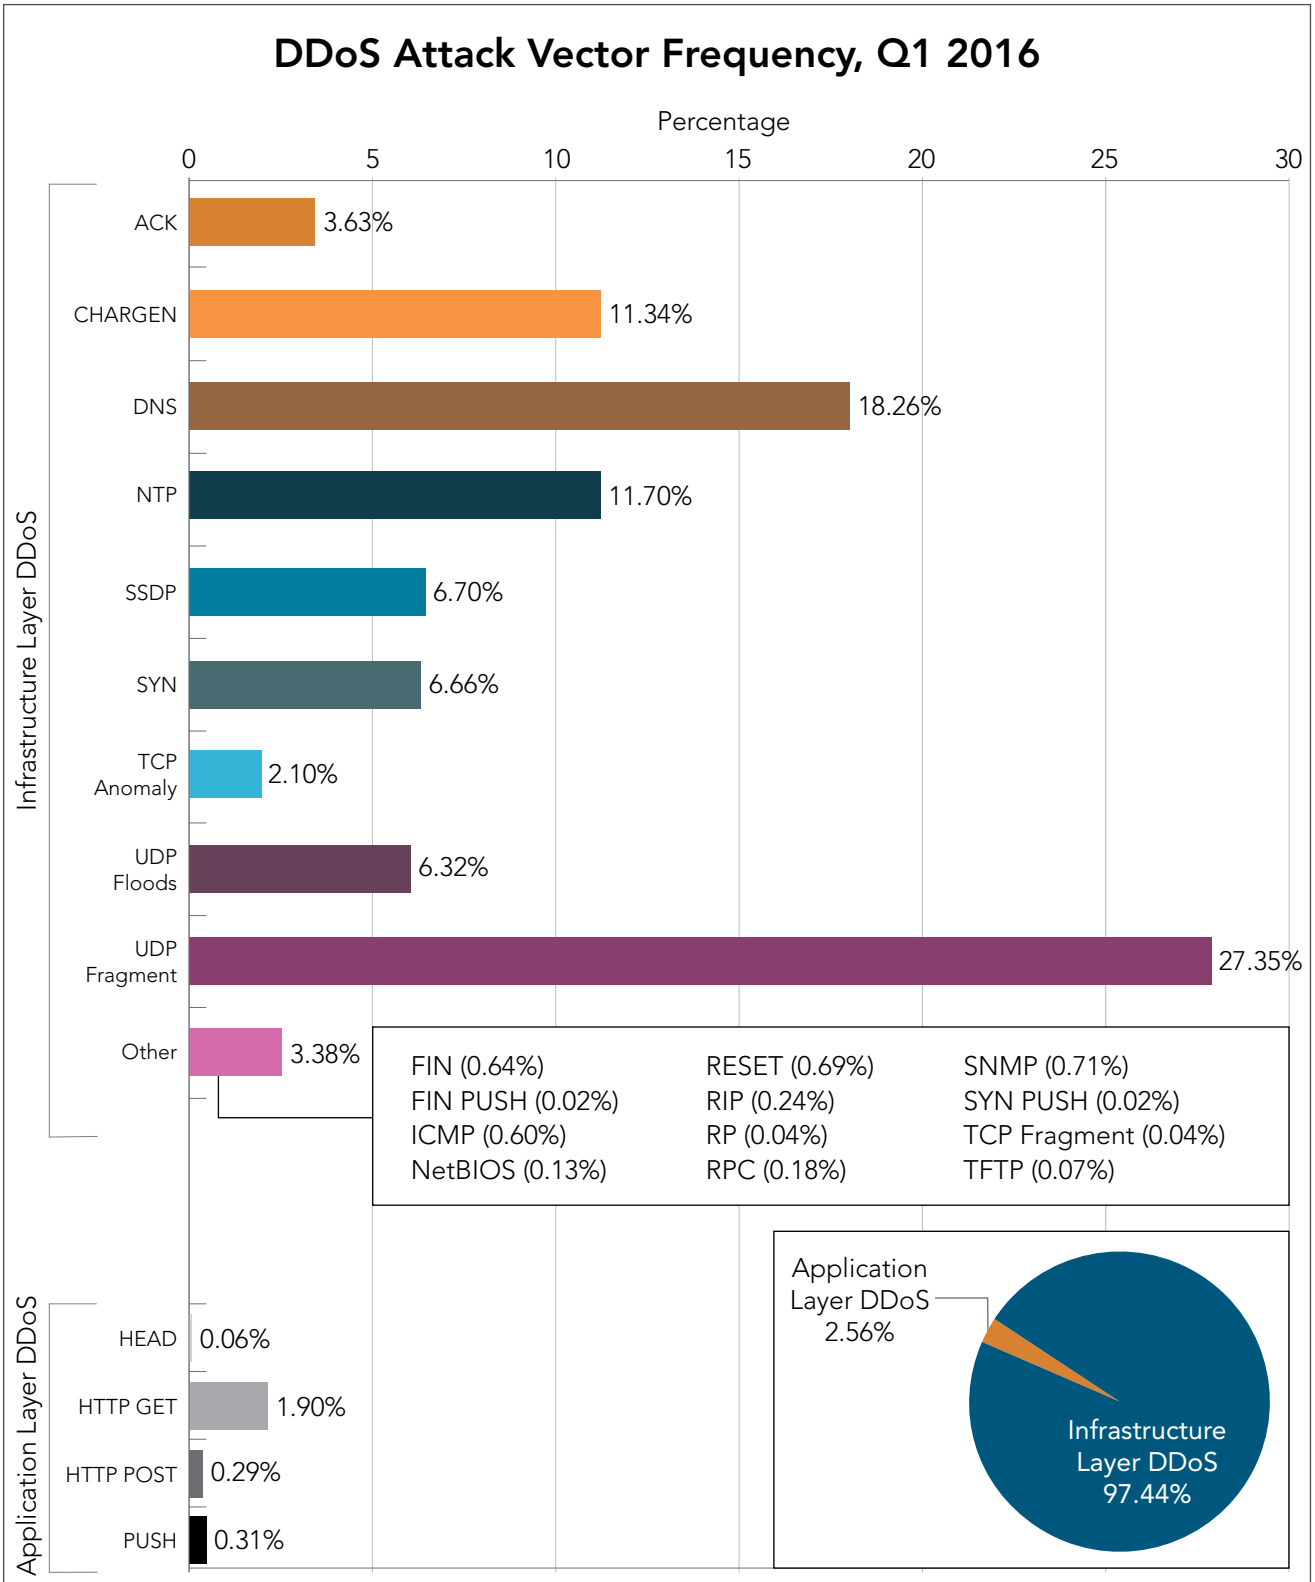

Figure 2-1: Four of the 25 DDoS attack vectors tracked this quarter—UDP Fragment, DNS, NTP, and CHARGEN—comprised nearly 70% of the attacks

At 2%, TCP anomaly attacks dropped a percentage point compared with Q4 2015. The TCP anomaly attack vector accounts for TCP floods that use uncommon, or anomalous, TCP flags in attacks. TCP anomaly attacks usually result from a combination of coding errors and intentional attack script choices.

In many cases, malicious actors modify well-known SYN flood scripts so that the flags in each packet are no longer just SYN flags. Some of these attacks don't have a SYN flag set any longer, yet otherwise have similar characteristics to SYN flood script attacks.

Documented coding errors in the Xor botnet TCP header assembly also resulted in attacks with up to three flag combinations.

When comparing the attack numbers from one year ago, we saw sharp increases in DNS, NTP, CHARGEN, and UDP fragment attacks. Conversely, SSDP and several application-layer DDoS attacks decreased. The decrease of SSDP and increase of NTP attacks in particular was surprising. Available SSDP reflectors — which typically deliver strong response amplification — decreased significantly in Q1. While a larger number of NTP servers were used as reflectors in Q1, as discussed in ***Section 4***, the number of NTP reflectors capable of providing the strong response amplification desired by malicious actors has decreased significantly. Still, attackers have increased the use of NTP over SSDP in the last few quarters.

Although we tracked two dozen attack vectors in Q1 2016, the top 10 vectors were responsible for the vast majority of the attacks.

To better understand the evolving threat landscape, we analyzed this subset of attack vectors during the past five quarters, as shown in Figure 2-2.

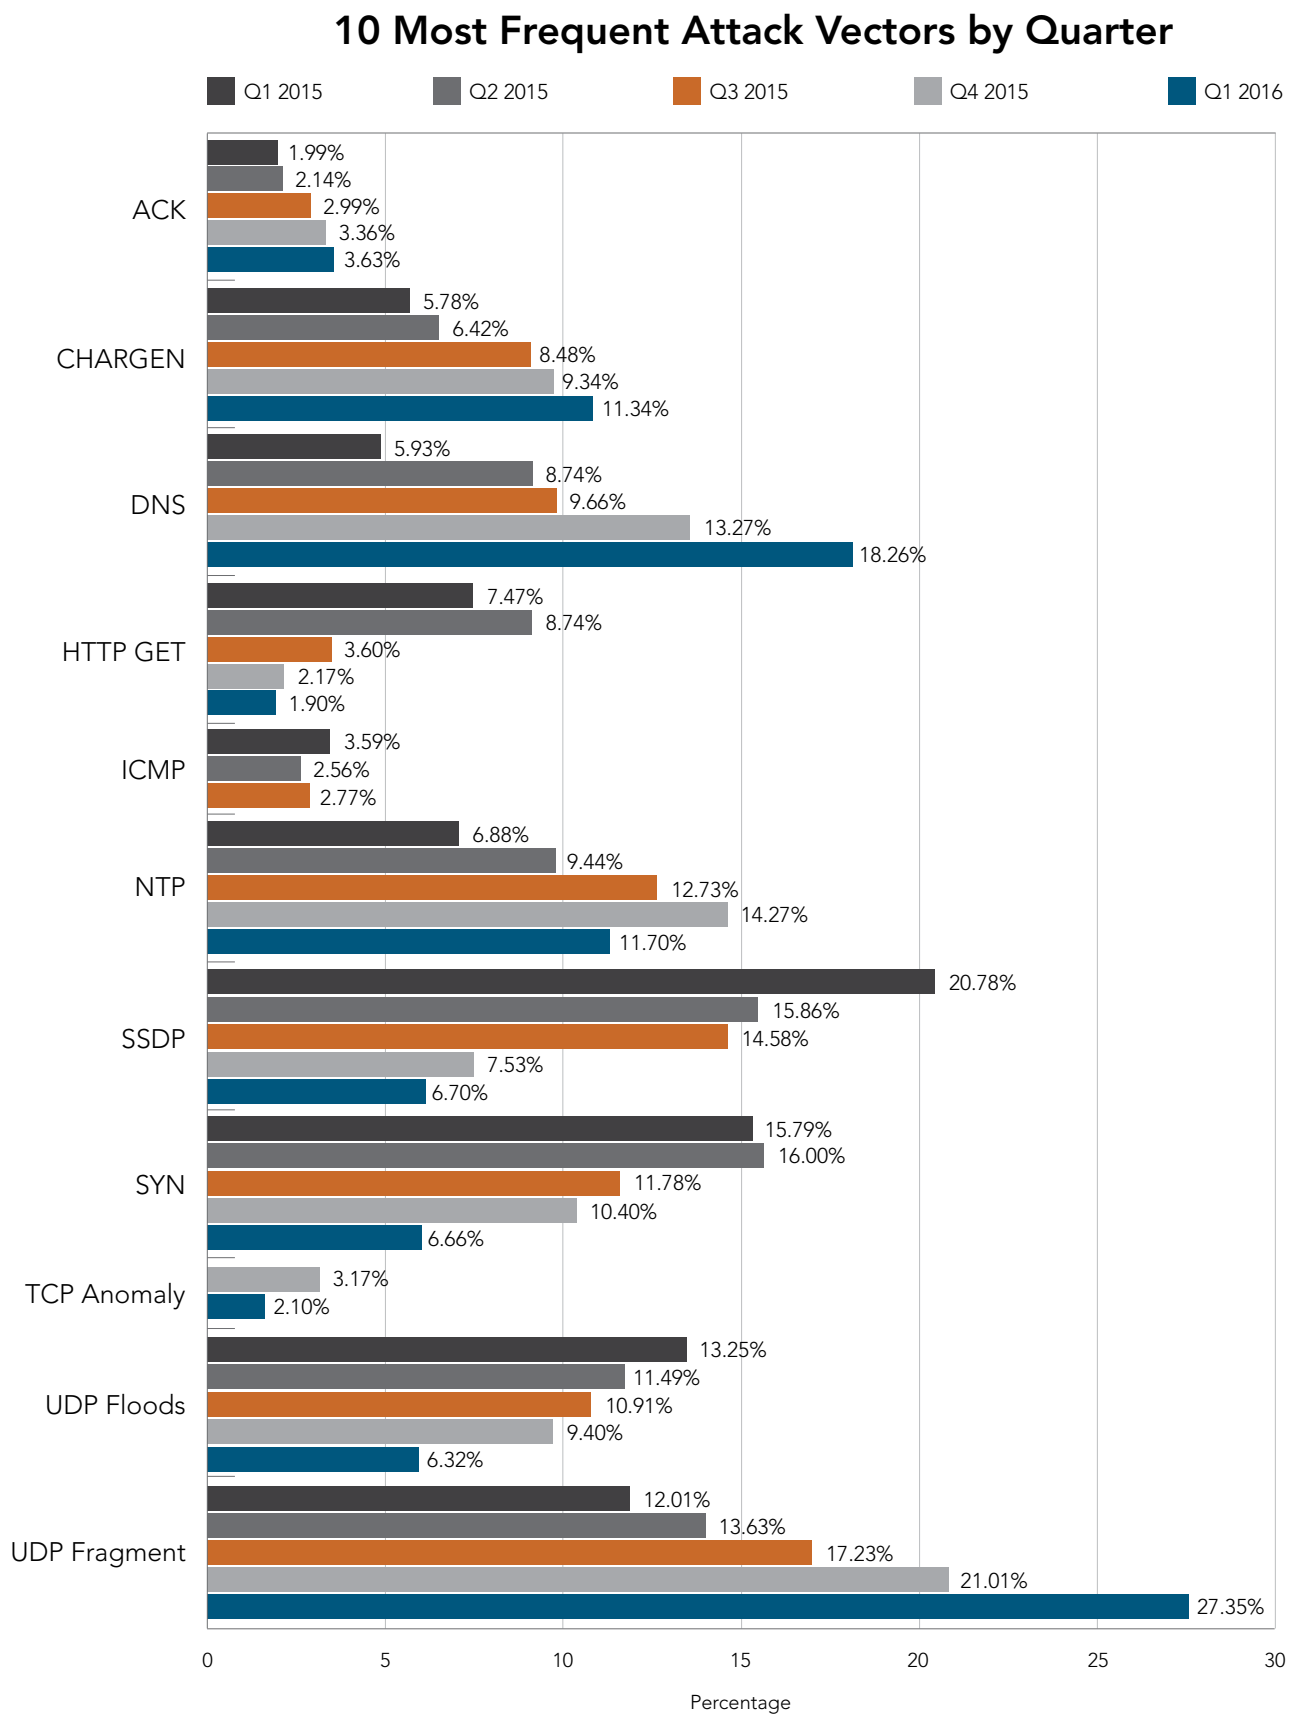

Figure 2-2: The 10 most popular attack vectors have remained consistent since Q1 2015, with the exception of TCP Anomaly attacks, which first edged out ICMP attacks in Q4 2015

For example, the reduction of SSDP attack traffic and increase in DNS and NTP attacks reflects the cyclical nature of attack tools and methods in the DDoS world. The overall trend shows a continued increase in infrastructure-based attacks, led by the use of reflection attack vectors. DNS in particular saw a large increase in attacks this quarter.

Starting this quarter, we modified the computation to include the complete dataset to determine vector percentages for each quarter, as opposed to including only the top 10 vectors. Therefore, the percentages in each column will not add up to 100% and will vary slightly from previous reports.

The continued rise of multi-vector attacks, as shown in Figure 2-3, suggests that attackers or their attack tools are growing more sophisticated. This causes problems for security practitioners, since each attack vector requires unique mitigation controls.

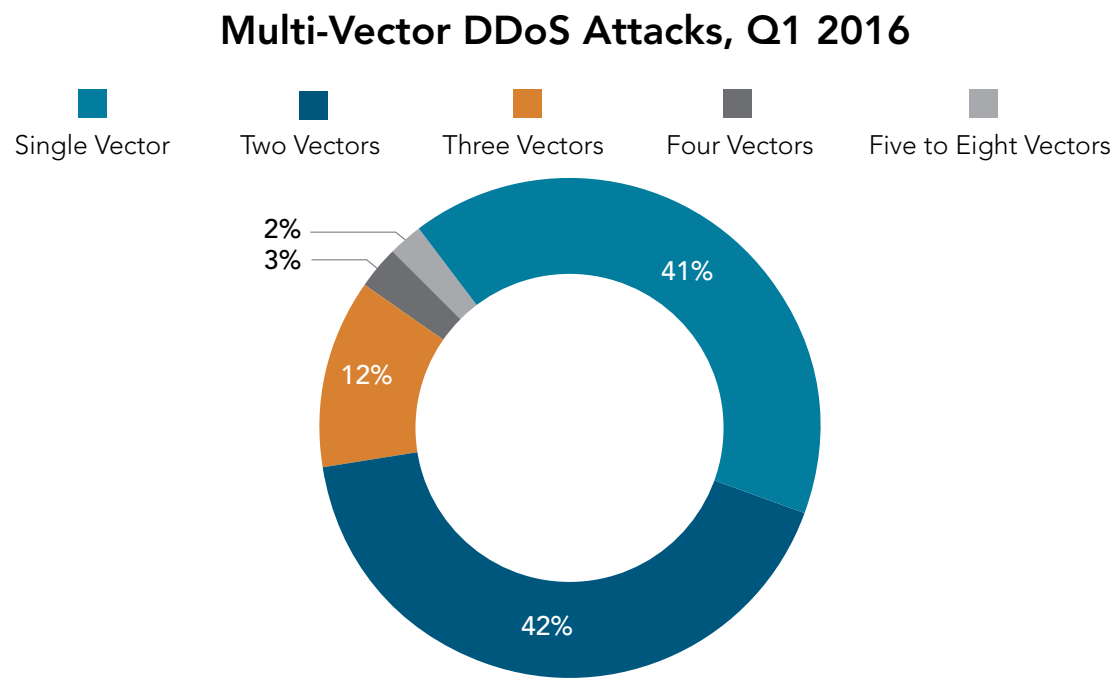

Figure 2-3: Multi-vector attacks accounted for 59% of DDoS activity in Q1 2016, reflecting a slight increase compared with last quarter (56%)

Booter sites have played a key role in enabling more multi-vector attacks. Many of the same attacks we identified throughout 2015 are included in these DDoS-for-hire frameworks, and multiple attacks can be launched simultaneously, depending on the service purchased. These sites offer an easy-to-use interface with a menu of attacks. Malicious actors may not even know which attacks are being launched due to the names used on these sites. For example, something labeled a UDP flood could actually be an NTP reflection flood or some other attack. The majority of attacks included in booter site frameworks are infrastructure-based (layers 3 and 4).

In all, 59% of all DDoS attacks mitigated in Q1 2016 were multi-vector attacks, building on a trend that has developed over the past year, as shown in Figure 2-4.

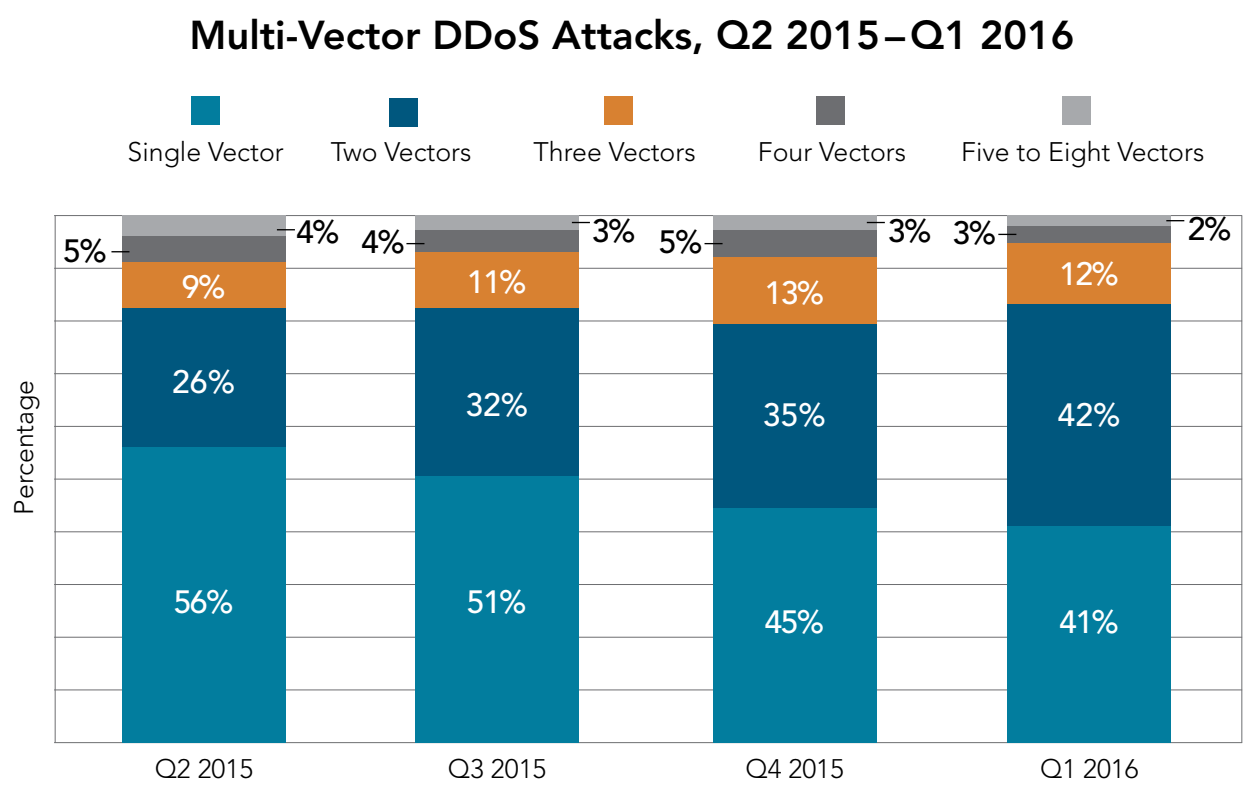

Figure: 2-4: The frequency of single vector attacks has steadily declined since Q2 2015, with the majority of attacks now using at least two attack vectors

2.2 / Mega Attacks / In Q1 2016, a record 19 attacks exceeded 100 Gbps, the largest hitting the software & technology, gaming, and media & entertainment sectors, as shown in Figure 2-5.

That's a significant increase from the five DDoS attacks that registered more than 100 Gbps in Q4 2015. Looking back further, there were eight mega attacks in Q3 2015 and 17 in Q3 2014 — the previous record.

In Q1 2016, the largest DDoS attack measured 289 Gbps, a 20-point drop over the largest attack in the previous quarter (309 Gbps).

Aside from that attack, most of the mega attacks seemed to use tools common to booters/stressers. Of the 19 mega attacks, the software & technology sector was targeted most frequently, including the 289 Gbps attack on Feb. 21. While new reflection methods are particularly hard hitting when first leveraged, the mitigation and cleanup of the associated vulnerabilities typically leads to a reduction in total attack bandwidth. The amplification value of NTP attacks, for example, has ramped down due to the `monlist` query being patched over time. Early attacks with DNS were more powerful as well. This latest wave of mega attacks included 14 campaigns with DNS reflection attacks. Although not as powerful as the early days of DNS reflection, these campaigns produced significant attack bandwidth partly due to the increased amplification factor produced by the use of Domain Name System Security Extensions (DNSSEC) domains. This latest wave of mega attacks includes 14 campaigns with DNS reflection attacks. More detail is provided in the ***DDoS attack spotlight section of this report***.

Six DDoS attacks exceeded 30 Mpps in Q1, and two attacks peaked at more than 50 Mpps, as shown in Figure 2-6. The packet rate affects some routers and networks more than the number of bytes because packets require more memory to track, tying up resources.

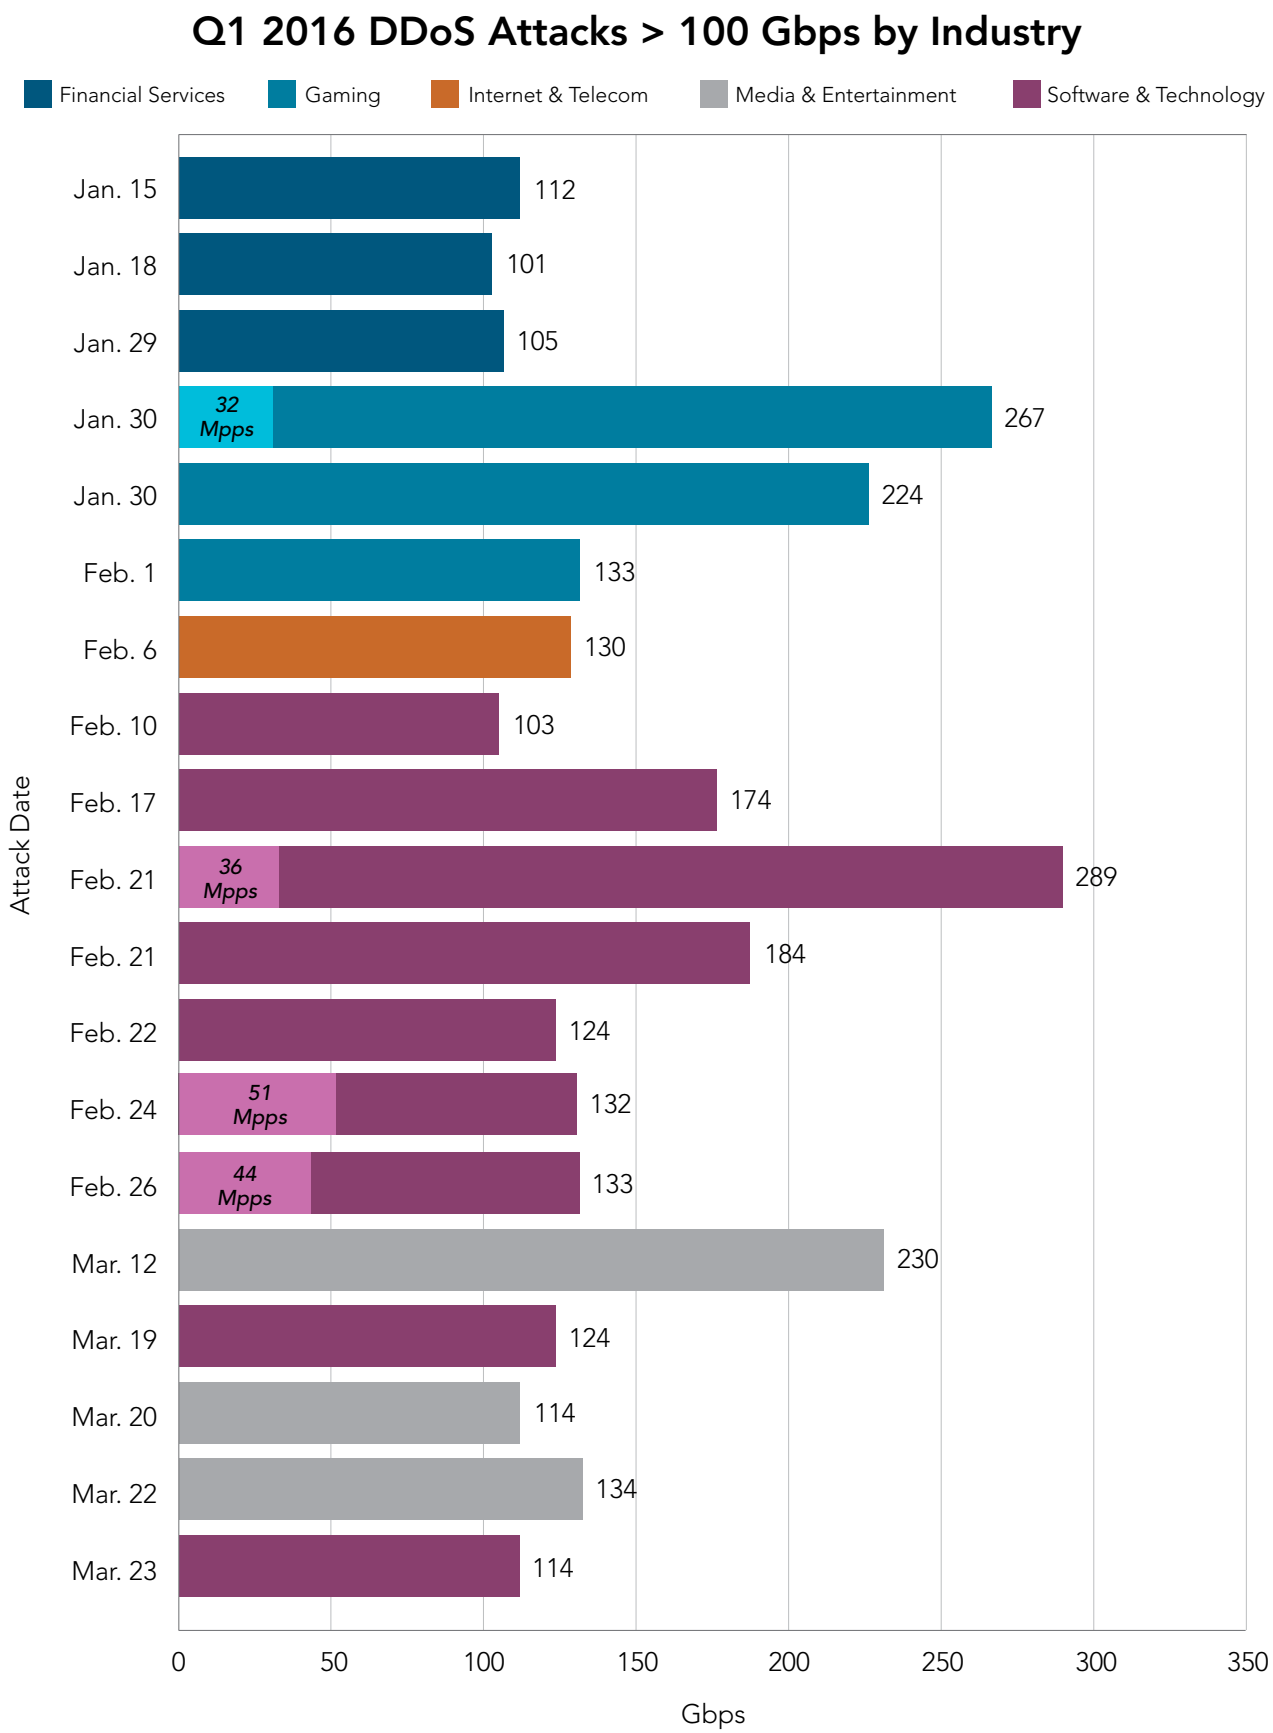

Figure 2-5: Nineteen attacks exceeded 100 Gbps in Q1 2016, with the largest hitting the software and technology, gaming and media-entertainment sectors

These high packet rates can result in packet loss within these routers and potentially cause collateral damage. The increased use of vectors with large packet size means fewer attacks should exceed this critical threshold. However, bandwidth exhaustion is still a concern.

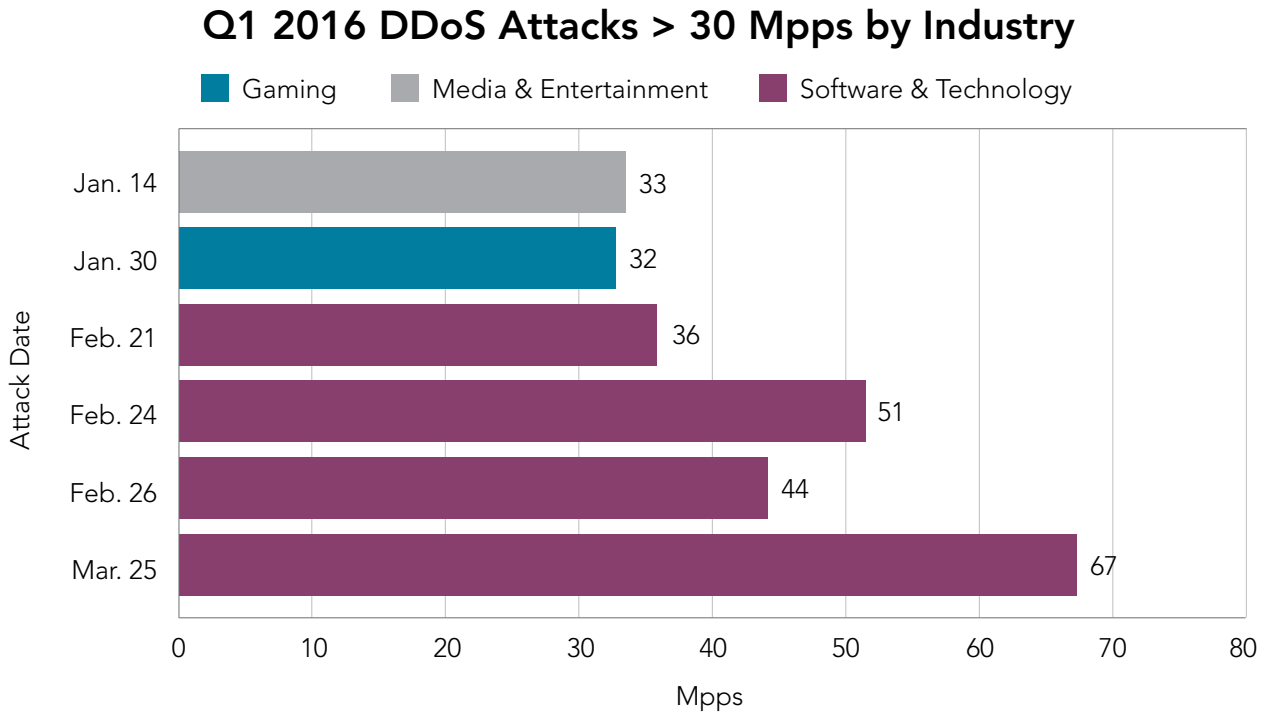

Figure 2-6: Of the six attacks exceeding 30 Mpps in Q1 2016, the four largest targeted the software and technology sector

2.3 / DDoS Attack Source Countries / As we've seen in many previous quarters, China was the top source country for DDoS attacks in Q1, accounting for 27% of all activity, as shown in Figure 2-7 and Figure 2-8. The US was the second-largest source of attacks (17%) while Turkey accounted for 10%.

Turkey's high placement on the list continues a trend we've seen in recent months where Eastern European countries are filling a void that had been occupied by Russia in recent years. One theory is that many attackers who had been operating from Russia have spread out to the Eastern European nations.

Top 10 Source Countries for DDoS Attacks, Q1 2016

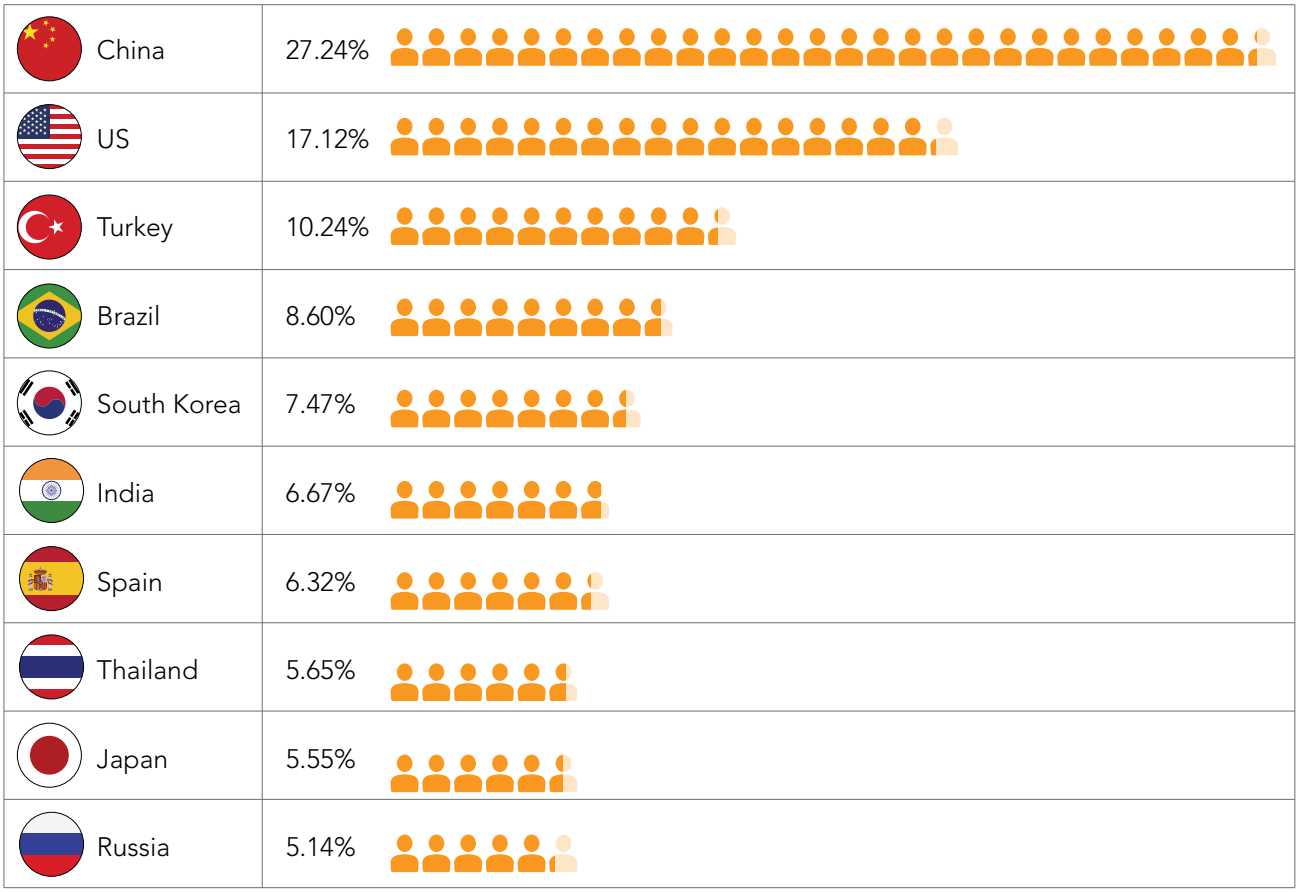

Figure 2-7: China was the top source of non-spoofed DDoS attacks in the first quarter, followed by the US

It is important to note that, while the packets we receive are from IP addresses in these countries, it is likely that a certain number of attackers used proxies or Virtual Private Servers (VPS) to hide their true originating IP addresses. This source information is further muddied by the fact that some of the available attack scripts rely on sending GET flood attacks from other websites, as is the case with pingback attacks, or are relayed through proxies.

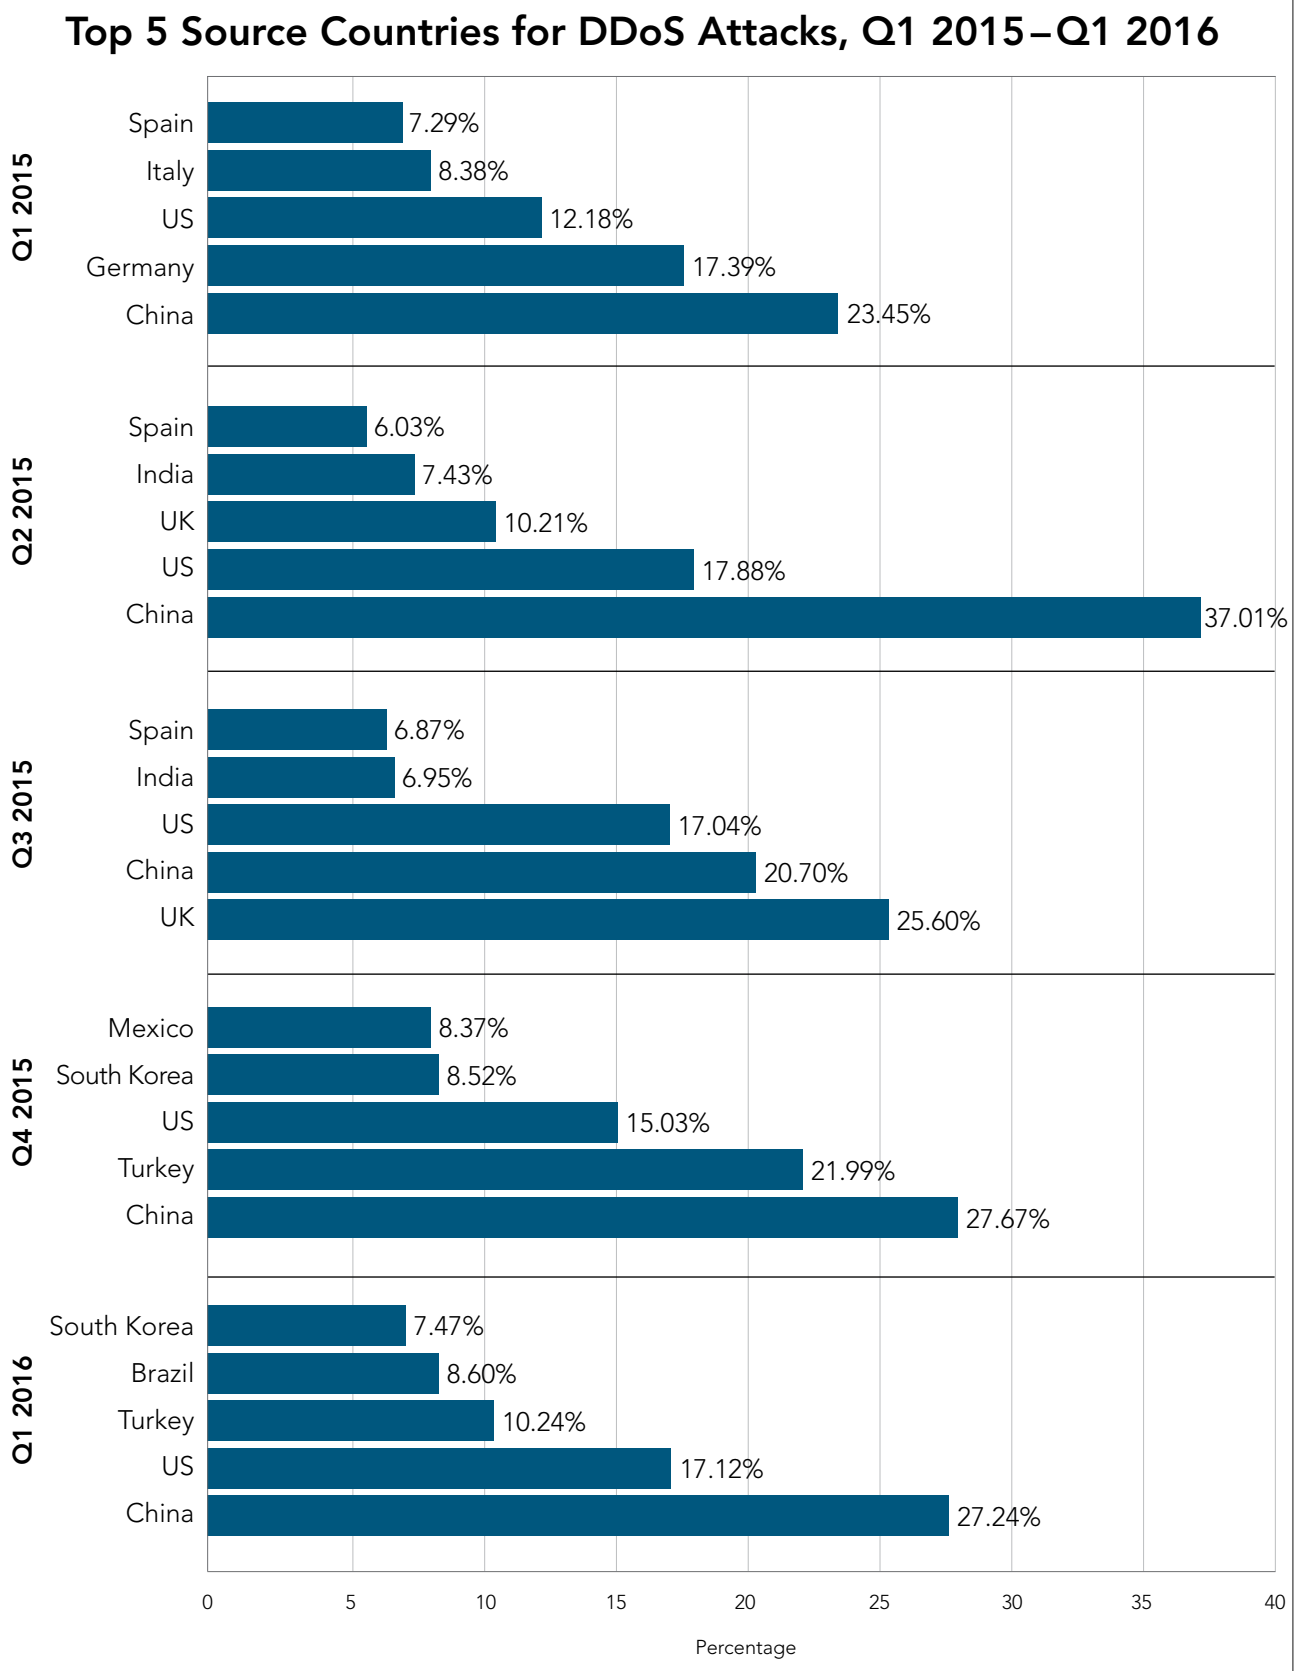

Figure 2-8: China has been the top source country for DDoS attacks since Q1 2015, with the exception of Q3 2015, when the UK took the top spot

2.4 / DDoS Attacks by Industry / The online gaming sector was hit hard in Q1 2016, accounting for 55% of all DDoS attacks, as shown in Figure 2-9. This was no surprise, since gaming has been the hardest-hit sector by DDoS attacks for the last few years. The frequency of attacks in that industry was about the same as in Q4 2015.

Gaming was followed by software & technology, which suffered 25% of all attacks in Q1, then media & entertainment (5%), financial services (4%), Internet & telecom (4%), education (3%), the public sector (2%), and retail & consumer goods (2%).

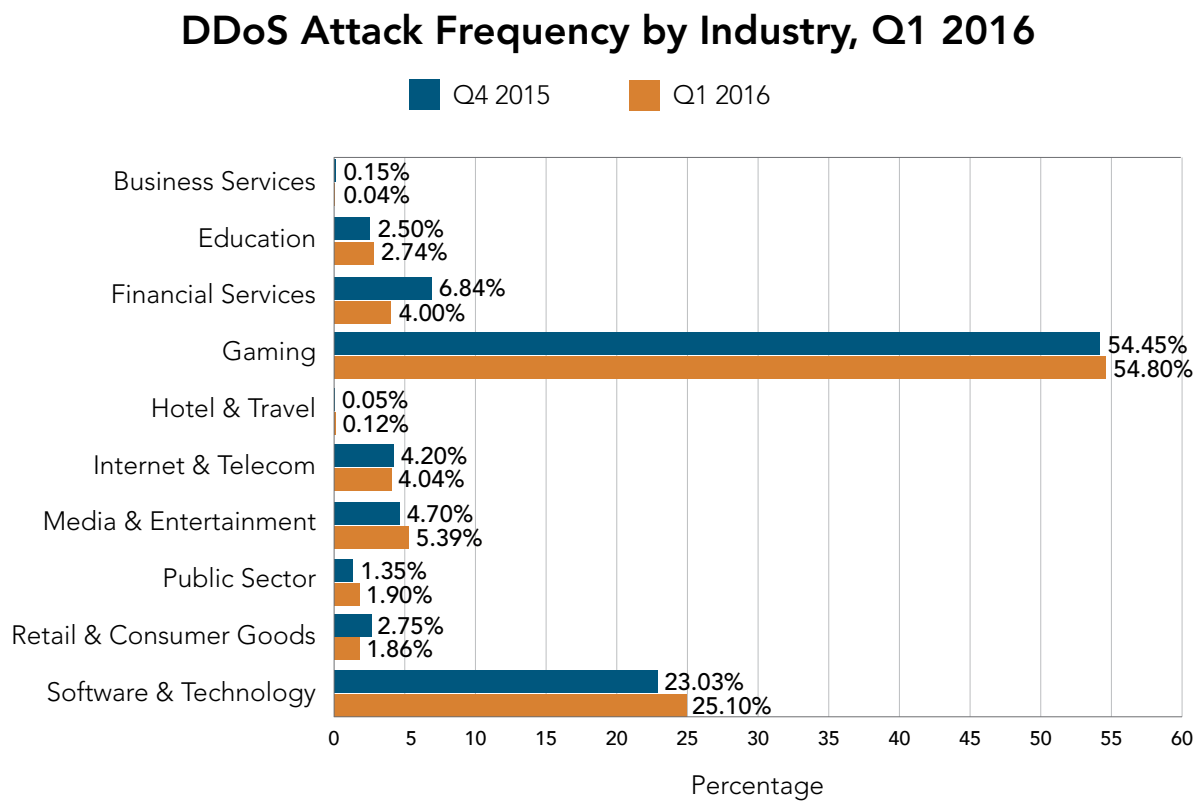

Figure 2-9: The gaming industry continued to be the most-frequently targeted sector for DDoS attacks, followed by the software and technology industry

Online gaming / Online gaming has remained the most targeted industry since Q1 2014, when we first began reporting on target industries. 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; to say nothing of disrupting competitive players.

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 that were launched for the holiday season. At the end of 2015, we saw a similar pattern.

As a target industry, in Q1 2016, online gaming has also followed the trend of more reflection-based 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. This has led to the commoditization of DDoS platforms.

These low-cost platforms allow malicious actors to launch DDoS attacks from anywhere in the world. For example, one platform recently advertised 7,200 seconds of DDoS attack traffic for US$69.99 — less than 2 cents per minute.

Software & technology / The software & technology industry includes companies that provide solutions such as SaaS and cloud-based technologies.

Internet & telecom / The Internet & telecom industry includes companies that offer Internet-related services such as ISPs and DNS providers. Attackers don't usually target an ISP directly. Instead, the attacks target sites hosted by a provider. The customer population of a hosting provider affects its likelihood of attack. Providers that host many sites or high value sites are at higher risk of 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 services industry includes major financial institutions such as banks, insurance companies, payment providers and trading platforms. Recently, the financial industry has been the focus of various extortion attempts and the Armada Collective led the way with multiple extortion attempts against financial services companies.

Media & entertainment / The media & entertainment industry saw about the same level of attacks in Q1 2016 as in Q4 2015: 5%.

2.5 / DDoS Attacks — A Two-Year Look Back / Long-term trend analysis shows that half of all attacks were between 400 Mbps and 5 Gbps in size.

While this is a considerable range, it's worth noting that there's a significant grouping of attacks just beyond the 5 Gbps threshold. Attacks between 3 and 10 Gbps account for more than 30% of all attacks.

While the mean attack size fluctuates significantly quarter over quarter, the median is much more stable and better represents what can be expected.

The size of documented attacks had steadily declined over the last year with very large attacks becoming less frequent while the median attack size has remained stable over time. Q1 2016 — with a record number of mega attacks — was an exception. This may be the start of a new trend or it may be a one-time anomaly. As we look at DDoS size and frequency as a function of time, as shown in Figure 2-10, one can't help but speculate on how long it will be before the right end of the chart will be represented as a more consistent, solid line.

If we look at the median attack size by quarter, Q2 2014 was the lowest at 1.16 Gbps. The highest mean attack size, 2 Gbps, was in Q1 2014 for the period shown. One factor contributing to the higher mean during that quarter was the use of NTP reflection attacks.

During the early days of NTP reflection attacks, the reflectable hosts responding to malicious monlist queries were plentiful. Today, more hosts have been patched for this vulnerability, which in turn has reduced this vector's impact. Q1 2014 also marked the first significant shift away from application layer DDoS attacks.

Application-layer DDoS attacks don't typically generate high bandwidth but can potentially cause significant degradation issues to a single target. The earlier quarters represented on this chart contained more application-layer attacks, which in turn resulted in lower mean attack size. The rest of the quarters have, for the most part, median values hovering around the 1.5 Gbps mark. Exceptions were Q1 2016 and Q4 2015, which each had a 1.8 Gbps median, and Q1 2015 with 1.3 Gbps.

While booter/stresser frameworks have evolved to more frequently support multi-vector attacks, the stable median size indicates they still have not evolved to support the larger-than-normal attack bandwidth available with traditional malware botnets.

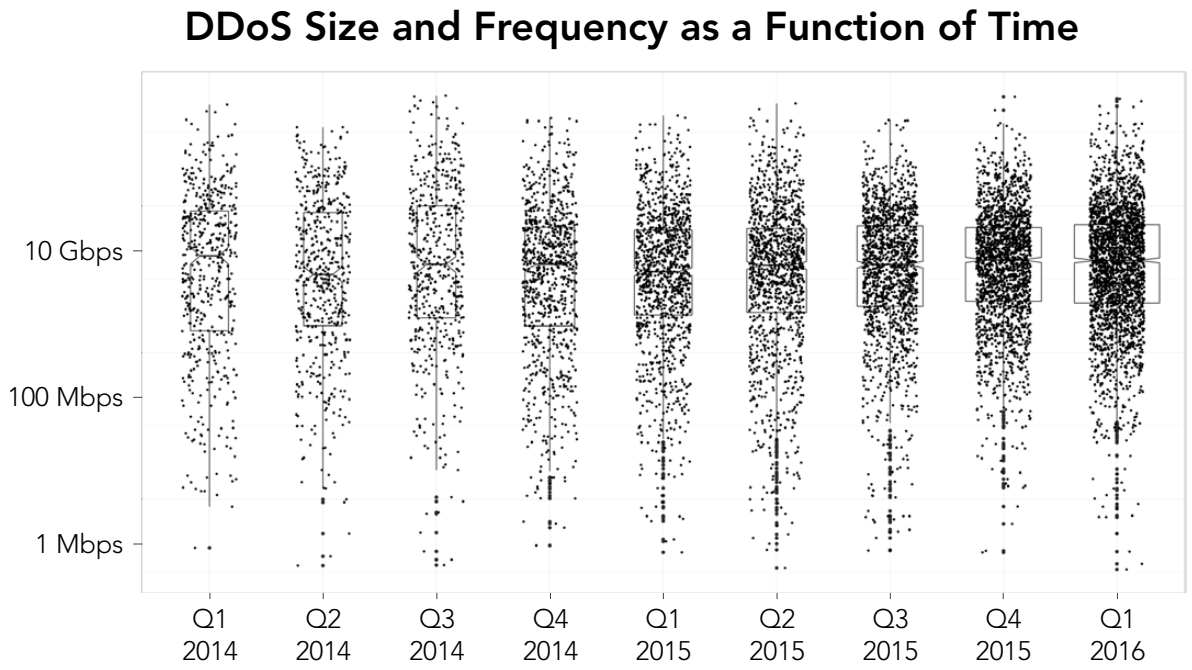

Figure 2-10: While the median size of DDoS attacks has varied only slightly in recent quarters, the number of attacks has grown dramatically

The boxes for each quarter represent the middle 50% of attacks by attack size, while each dot represents an individual attack. The vertical axis has a logarithmic scale; the upper attacks are many thousands of times larger than the bottom ones.

The median packet rate has remained under the 1 Mpps mark for the past two years, as shown in Figure 2-11. More recently, in the past three quarters, the median packet rate has remained under 0.5 Mpps.

In theory, a 1 Gbps interface should be able to send a packet volume exceeding 1 Mpps. Still, there are factors limiting a single host, such as the bandwidth available from the ISP or even congestion points in the path to their target.

The few attacks exceeding 200 Mpps in Q3 and Q4 of 2015 and earlier quarters were an exception. These were indicators of large DDoS botnets and well-connected, powerful servers. These high packet rates would likely hinder or completely halt communications on low to even mid-range networking devices.

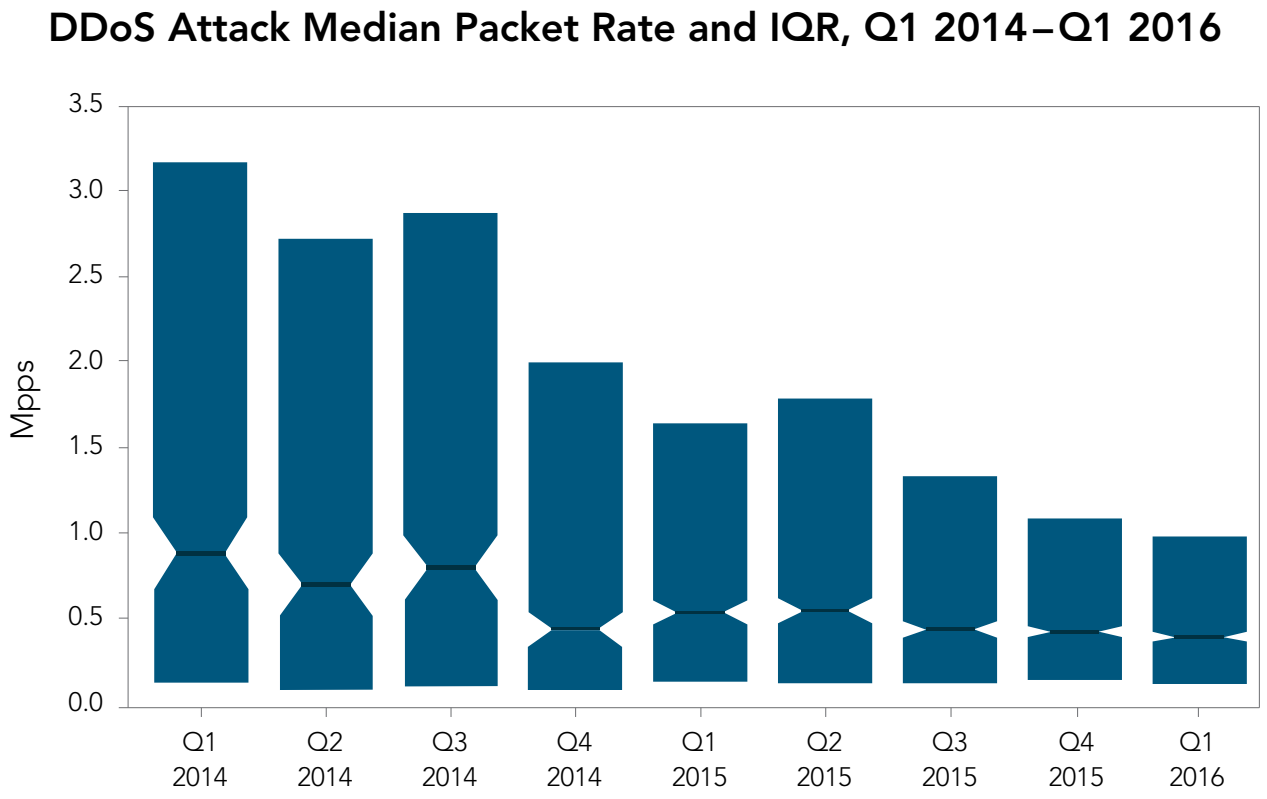

Figure 2-11: While there were six attacks that exceeded 30 Mpps (not shown), more than half of the DDoS attacks in Q1 2016 measured 1 Mpps or less

The graph shows the packet rate for the middle 50% of DDoS attacks from Q1 2014–Q1 2016. The top and bottom 25% are excluded as outliers.

So far, this kind of DDoS power has not been readily obtainable. However, with constantly evolving DDoS malware, high packet rate attacks are something that must be considered for DDoS mitigation.

2.6 / Reflection DDoS Attacks, Q1 2015 – Q1 2016 / The Sankey graphic in Figure 2-12 shows how DDoS reflection attacks have trended during the past five quarters. Through the routed network, we tracked 10 infrastructure-layer DDoS-reflection vectors. The most used vectors seem to 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 are the most used reflection DDoS vectors. DNS has now surpassed SSDP as the top vector. The use of SSDP reflection was most prevalent in Q1 2015 followed by Q3 2015, from there it has dropped off. DNS has continued a steady increase in use with a peak in Q1 2016.

SSDP, NTP, DNS, and CHARGEN have consistently been used as the most common reflection attack vectors, as can be seen on the left axis, and the use of reflection attacks has increased dramatically since Q1 2015, as shown on the right axis.

Internet of Things (IoT) devices are coming to market faster than is sustainable from a security perspective, which may be driving the changes seen here. Many IoT devices, such as printers, are shipped with little or no due diligence from a security perspective. These devices are home-based and cannot be effectively updated or managed by the end user. As a result, these exposed services can potentially be incorporated into distributed attack platforms.

Looking on the right, from top to bottom, shows a steady increase in the use of reflection-based DDoS attacks each quarter. A big takeaway from the Sankey graph is that malicious actors are finding it more profitable to choose reflection over

Reflection-Based DDoS Attacks, Q1 2015–Q1 2016

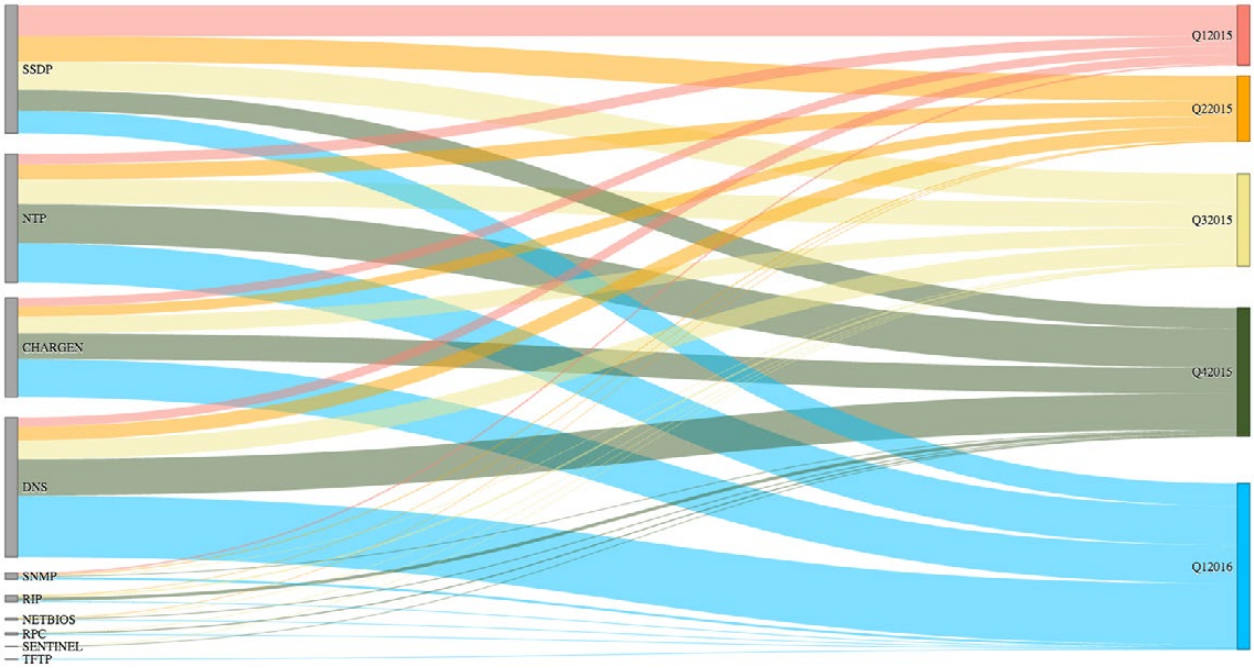

Figure 2-12: SSDP, NTP, DNS, and CHARGEN have consistently been used as the most common reflection attack vectors, as shown on the left axis. The use of reflection attacks has increased dramatically since Q1 2015, as shown on the right axis.

infection. Instead of spending time and effort to build and maintain DDoS botnets, it is far easier for attackers to trick existing network devices into participating in attacks. This shift in methodology has spread to the DDoS-for-hire ecosystem.

The percentage growth in reflected DDoS attacks over the past five quarters is shown in Figure 2-13.

Distribution of Total DDoS Reflection Attacks, Q1 2015–Q1 2016

Quarter	Percentage
Q1 2015	11.78%
Q2 2015	12.86%
Q3 2015	18.01%
Q4 2015	25.01%
Q1 2016	32.33%

Figure 2-13: The frequency of reflection-based DDoS attacks has grown significantly quarter over quarter

2.7 / Repeat DDoS Attacks by Target / A new trend we've started exploring is the number of repeat attacks against the same organization.

There were an average of 15 attack events per targeted customer in Q1 2015. By Q4 2015, the average attack events per target rose to 24. And in Q1 2016, the average number of attacks grew to 29, as shown in Figure 2-14.

In the past, many attackers would see that a site or network was protected and move on. Now, the latest trend is for attackers to hammer away at high-value organizations, regardless of effect, looking for a moment when defenses might drop. This is frequently the case with gaming organizations, where even the slightest latency can have a noticeable effect on online gamers.

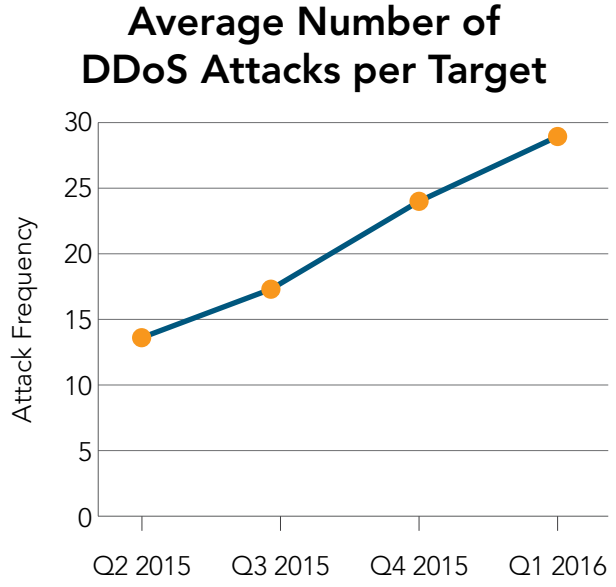

Figure 2-14: In Q1 2016 there were an average of 29 DDoS attacks per target, up from 24 last quarter

Another reason for the continued increase in repeat attacks is that acquiring and outfitting DDoS attack platforms has become cheap and easy to use. DDoS attacks have also been used as a diversion technique to exhaust company resources while attacks are launched against the primary target. For example, a DDoS attack may be used to hide the attackers' true purpose: data exfiltration.

The most frequently attacked customer in Q1 2016 was targeted with 283 DDoS attacks. That represents three separate attacks per day. Attackers seem to be becoming more persistent in trying to disrupt services against their targets.

2.8 / DDoS Attack Spotlight / During Q1 2016, an alarming number of attacks — 19 — exceeded 100 Gbps. But what was even more disturbing were the methods used in these attacks. Many of these attacks had characteristics similar to common booter/stresser attacks. It appears that it is now common for what appear to be booter/stresser-based attacks to create floods of 5 to 50 Gbps. The increase of the output bandwidth of these lowest-common-denominator tools puts ever-larger attacks in the hands of unsophisticated attackers.

This quarter, many attacks with a similar style to the common booter-based attacks exceeded 100 Gbps. Our spotlight attack this quarter is one such attack. It peaked at 230 Gbps, using only a DNS reflection and amplification attack vector. In fact, 14 of the 19 mega attacks contained at least some level of DNS reflection.

Media lash out / The target of the spotlight attack was an Akamai customer in the media industry. Media sites often become targets based on the news they publish. Groups such as Anonymous often use DDoS attacks as a form of digital protest.

The tools used in these attacks are accessible to even lesser-skilled actors. Using sites offering stress-testing services makes it difficult for the target to determine true attribution.

The spotlight attack was distributed and mitigated across five scrubbing centers, as depicted in Figure 2-15. The bulk of the traffic was routed through Akamai's Frankfurt scrubbing center.

The attack included both UDP fragment and DNS attack vectors. The attack signature is depicted in Figure 2-16. One standout shown in the signatures is the use of domains that were DNSSEC-enabled. This is not a new attack method; however, the increased use of DNSSEC enabled legitimate domains is being leveraged by malicious actors in the latest flurry of attacks.

Akamai SIRT published a ***security bulletin*** on the increased use of DNSSEC in DNS DDoS attacks in February 2016.

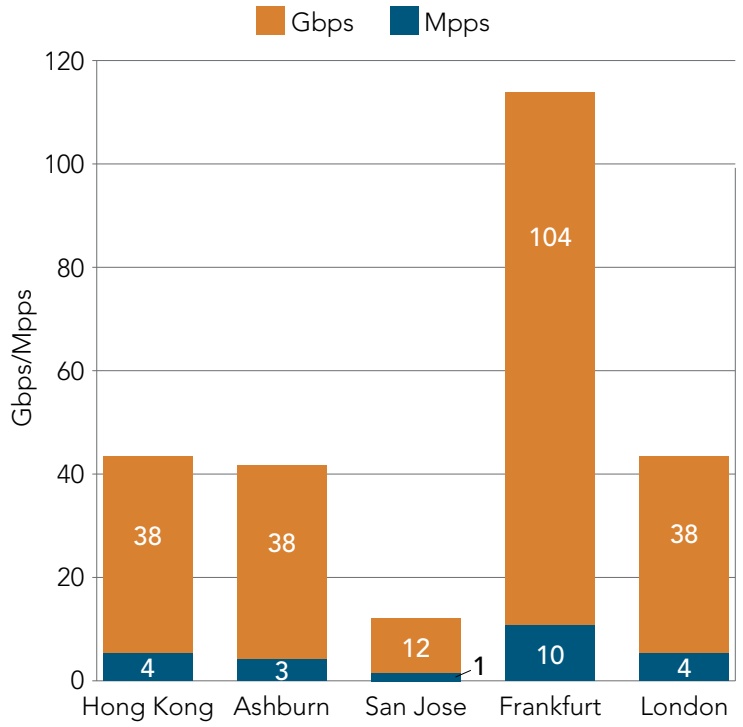

Figure 2-15: Attack traffic distribution within scrubbing center locations, highlighted with Frankfurt absorbing the highest peak bandwidth at 104 Gbps

In the earlier days of DNS attacks, it was common to see domains that were purposely set up to produce the highest possible amplification factor. By using a legitimate domain with DNSSEC-enabled, mitigation becomes a bit more complicated.

If all DNS resolution requests for these domains are dropped, this could potentially result in a Denial of Service (DoS) effect for legitimate users of the abused domain. In addition to mitigation issues for the DNS resolver, some DNSSEC-enabled domains have produced response sizes in excess of 4,000 bytes. This extra-large payload size makes DNSSEC-enabled domains attractive to use in reflection attacks.

The problem is that open DNS resolvers, spread throughout the Internet, will attempt to resolve any domain, even if the server does not have the domain configured locally. In the latter case, the DNS resolvers will send a query out to the domain and will often cache the response. During an attack, a rapid rate of queries is continuously sent to the open DNS resolver, which dutifully responds to the listed source address.

```
Attack Signature:
23:35:05.205576 IP A.A.A.A.53 > B.B.B.B.4444: 12545| 22/0/0 TXT "v=spf1
ip4:X.X.X.X ip4:X.X.X.X ip4:X.X.X.X mx a:<redacted> -all", A X.X.X.X,
AAAA X.X.X.X, DNSKEY, DNSKEY, DNSKEY, DNSKEY, Type51, RRSIG[|domain]

23:35:05.205587 IP A.A.A.A.53 > B.B.B.B.4444: 31577| 22/0/0 AAAA
X.X.X.X, DNSKEY, DNSKEY, RRSIG, RRSIG, RRSIG[|domain]

23:35:05.205600 IP A.A.A.A.53 > B.B.B.B.4444: 20265| 20/0/1 MX <re-
dacted>. 5, MX <redacted>. 5, TXT "v=spf1 ip4:X.X.X.X ip4:X.X.X.X ip4:X-
.X.X.X mx a:<redacted> -all", A X.X.X.X, AAAA X.X.X.X, DNSKEY, DNSKEY,
DNSKEY, DNSKEY, Type51, RRSIG[|domain]

23:35:05.205632 IP A.A.A.A.53 > B.B.B.B.4444: 31577| 22/0/0 RRSIG, RR-
SIG, RRSIG, RRSIG, RRSIG[|domain]

Fragmented UDP packets due to large response size:
23:35:05.205663 IP A.A.A.A.54 > B.B.B.B: udp
23:35:05.205695 IP A.A.A.A.26 > B.B.B.B: udp

Reply from unresponsive DNS server(Server no longer listening for DNS
queries on port 53)
23:35:04.874589 IP A.A.A.A.45 > B.B.B.B: ICMP A.A.A.A.45 udp port 53 un-
reachable, length 73

Reply from target IP, port is unreachable(The target does not listen on
UDP port 4444)
23:35:04.874749 IP B.B.B.B > A.A.A.A: ICMP B.B.B.B udp port 4444 un-
reachable, length 36
```

Figure 2-16: Attack signatures revealing the use of a DNSSEC enabled domain

When the recursive resolvers use DNS caching, many response packets will be sent without causing additional stress to the authoritative domain server that actually hosts the DNS record for the reflected domain.

DNS queries with a similar structure to the attack script can be reproduced using dig on a Linux command line. Shorter domain names help achieve higher potential amplification factors since the query size is reduced. We observed attacks in Q1 2016 with query payloads in the range of 35 – 50 bytes, depending on the domain used. Figure 2-17 shows the options observed in the DNS queries.

```
Options set in DNS query:
*Force reply in UDP
*UDP payload size fixed: common values include 9000,65000, and 65527
*Source of the query is spoofed with the IP of the intended target(Just
like any other UDP based reflection attack)
*Recursion desired
*ANY record query type
*OPT pseudo-RR
```

Figure 2-17: Attack tool DNS query recipe and command-line equivalent

Since the attackers send out UDP queries, the DNS resolvers replied using UDP. Using `ignore` on the command line prevents repeat TCP retry authentication requests by the client once a truncated response (more than 512 bytes) is detected.

In an attack scenario, the target host receiving the unsolicited DNS reply traffic would not attempt a TCP retry, since it never made the request. In reality, the traffic would likely never reach the target server due to filtering at the edge of the network. But it can still cause issues due to the amount of traffic generated by the large DNS response floods. In this case, the DNS flood generated 230 Gbps.

Using `OPT pseudo-RR` supports the use of Extensions Methods for DNS (EDNS), which in turn allows attackers to set the payload size. In the attacks we observed, the malicious queries were often set to 9,000. We have also observed payload sizes of 65,000 and 65,527.

The queries were made recursively. This means the open DNS resolver did not need to be authoritative for the queried domain. Instead, the open DNS resolver would query other DNS servers to resolve it. Eventually the response would be cached, and subsequent queries received responses from the open DNS resolver.

[SECTION]3
WEB APPLICATION ATTACK ACTIVITY

Akamai's Threat Research Team concentrated its 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 is to look at some of these common web app attack vectors, and identify the characteristics of the attacks as they transit our global network.

Additionally, in order to provide value to our readers, Akamai's Threat Research Team has filtered out traffic from third party commercial web vulnerability scanning vendors, which are often used for compliance testing. This traffic does not constitute real attack data and artificially inflates raw numbers.

3.1 / Web Application Attack Vectors / In Q2 2015, we added two attack types to the web application attacks we analyzed: XSS and Shellshock.

Including events based on Shellshock nearly doubled the number of attack events we analyzed in Q2 2015 vs. Q1 2015, with 173 million Shellshock attacks against Akamai customers in that one quarter. The Shellshock vulnerability was first discovered in September 2014 and received heavy media attention. As a result, this vulnerability has been patched on many systems. We expect the number of attempts to exploit it to continue to drop. We see evidence of this in Figure 3-1, where Shellshock only accounts for about 6% of attacks over HTTP. SQLi and LFI were the most frequent attack vectors over HTTP in Q1 at 47% and 35%, respectively.

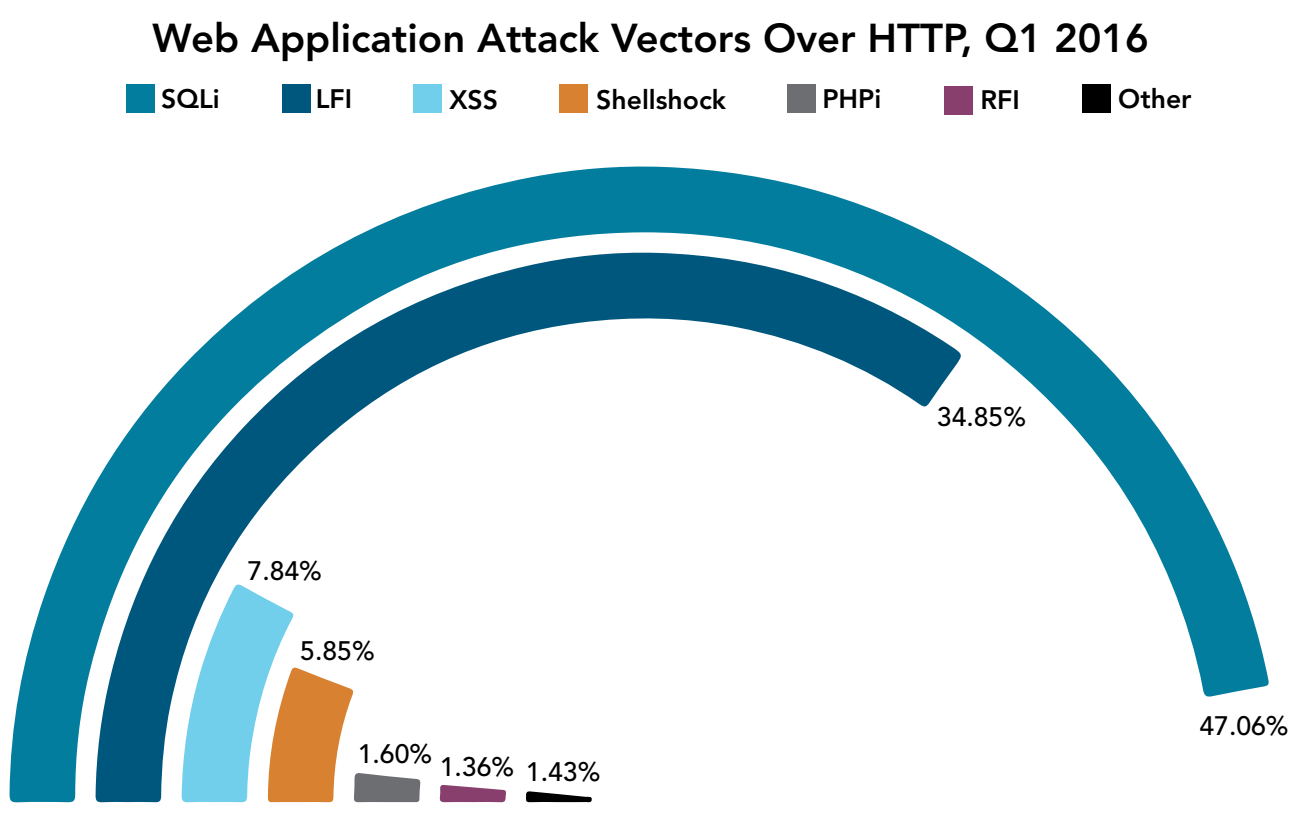

Figure 3-1: The three most popular attack vectors—SQLi, LFI, and XSS—were used in 90% of the attacks over HTTP

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.

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.

For web application attack vectors over HTTPS, LFI and SQLi accounted for 38% and 31% of the attacks in Q1 2016, respectively, as shown in Figure 3-2.

These attack vectors have remained effective due to coding issues that continue to crop up in websites. SQLi, which has been on the Open Web Application Security Project (OWASP) top 10 list for the better part of a decade, is a problem that can be solved through coding techniques that include security checks. In the rush to move applications to the web layer, however, many companies still do not check for SQLi vulnerabilities and thus leave themselves vulnerable to attack.

Over HTTPS, we saw Shellshock (20%) was used much more than it was over HTTP.

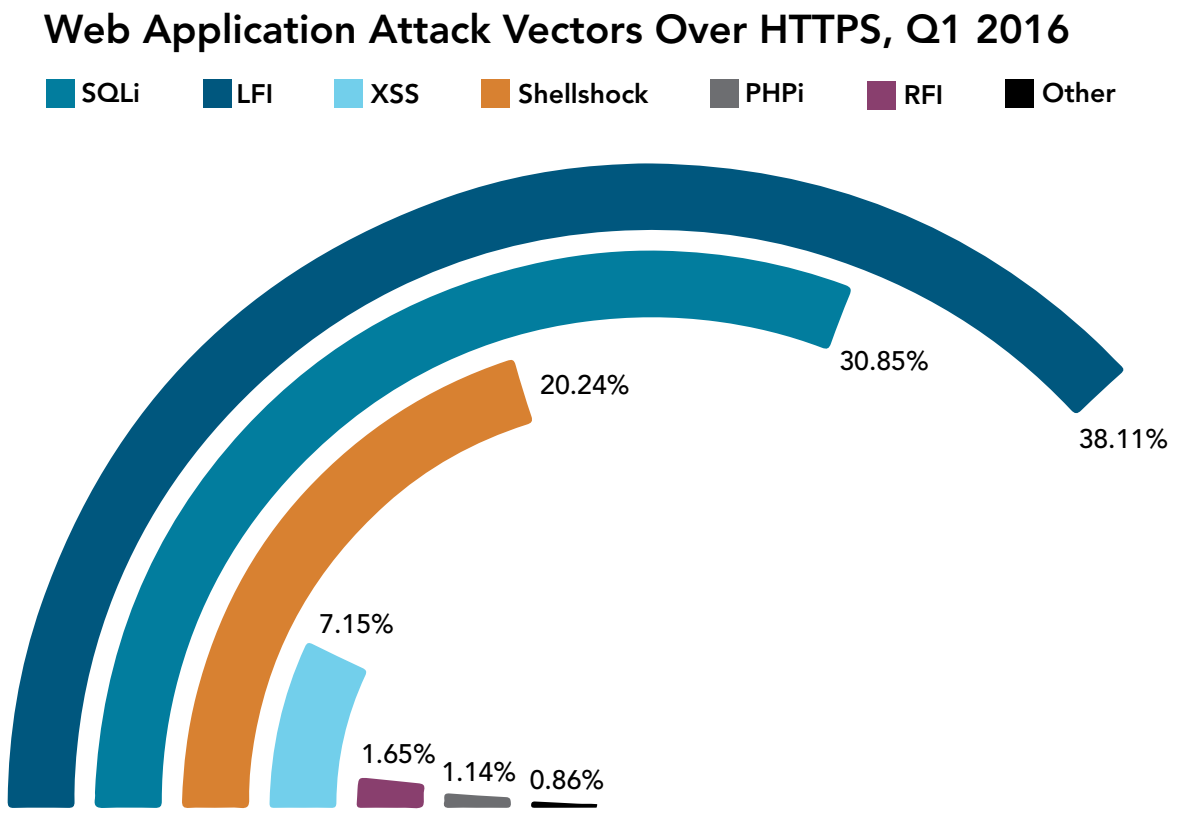

Figure 3-2: LFI, SQLi, and Shellshock were used in 89% of the attacks over HTTPS during Q1 2016

The majority of attacks — 70% — came over unencrypted channels (HTTP). This dominance has remained constant for the past year. The remaining 30% came over HTTPS, as shown in Figure 3-3.

A large percentage of websites either don't use HTTPS for their web traffic or use it only to safeguard certain sensitive transactions (such as login requests). However, HTTPS-based attacks still account for millions of attack alerts each quarter. Changes in how browsers represent the safety of HTTP sites will likely drive increase adoption of HTTPS. Time will tell how quickly the attackers adapt to this change.

Encrypting connections over HTTPS only affords protection to the data in flight. It does not provide any protection mechanisms for web applications, as attackers tend to shift to HTTPS to follow through on vulnerable applications.

Web Application Attacks Over HTTP vs. HTTPS, Q1 2016

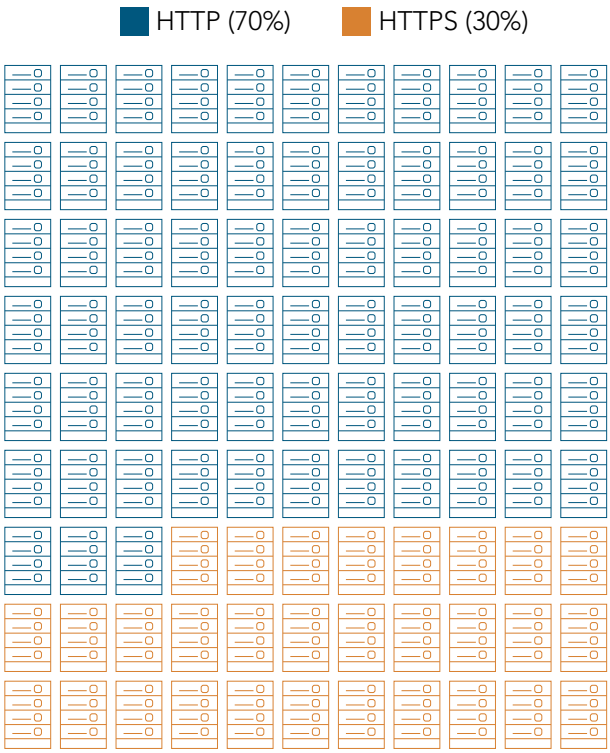

Figure 3-3: 30% of the web application attacks observed in Q1 2016 were over encrypted (HTTPS) connections, an increase from only 11% the previous quarter

3.2 / Top 10 Source and Target Countries / In Q1 2016, the US was the main source of web application attacks, accounting for 43% of attack origin traffic, as shown in Figure 3-4. This was a 13% drop from last quarter. Brazil was the second-largest source country at 12%, followed by China and the Netherlands (8% each), Romania (7%), Russia and the UK (6% each), Germany (4%), and the Ukraine and India (3% each).

The web application attacks we analyzed occurred after a TCP session was established. 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. Methods to obscure the source of these attacks include the use of proxy servers, rather than the direct packet-level source address manipulation

Top 10 Source Countries for Web Application Attacks, Q1 2016

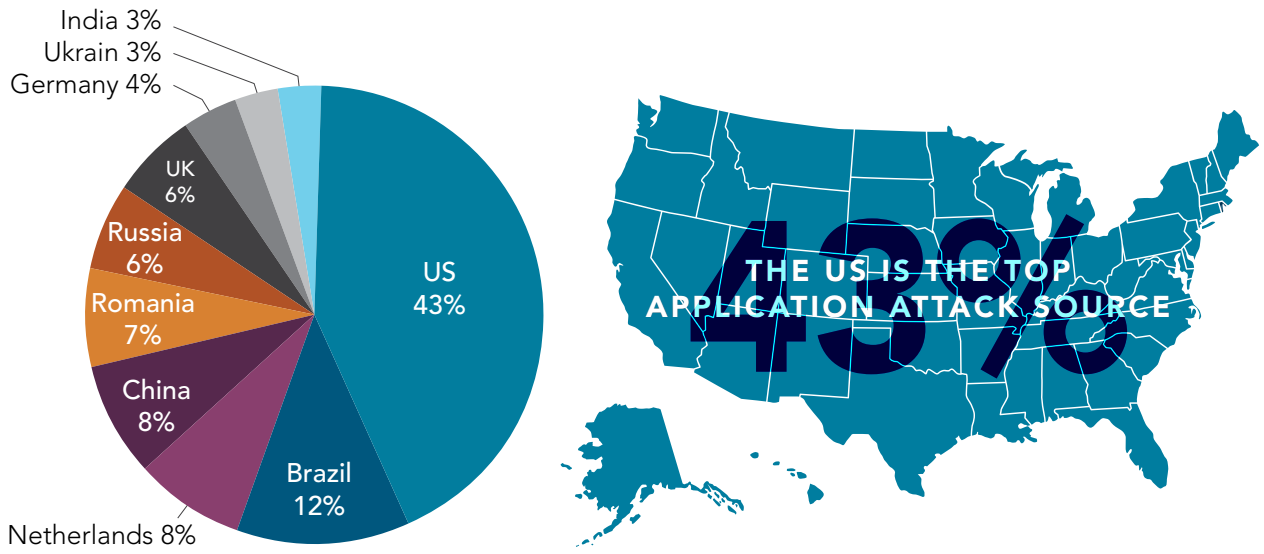

Figure 3-4: While the volume of attacks originating in the US dropped from 56% last quarter to 43% this quarter, the US remained the top source of web application attacks

commonly seen in the UDP-based infrastructure attacks. In contrast, when the attack source was the US, the main attack targets were in the retail industry, followed by manufacturing and media. In those cases, the most common attack methods were SQLi, LFI, and RFI.

In recent months, a global and respectable cloud IaaS provider opened data centers in Brazil. Since the opening of the data centers, Akamai has seen a large increase in the amount of malicious traffic coming out of Brazil, and specifically from the aforementioned data centers. Many of those attacks were against a Brazilian customer in the retail industry. A big difference with attack sources from Brazil was that the main targets were not only the US, but also India and Australia. When the attacks originated in Russia, the destinations were mostly in the retail industry in the US and the UK.

Countries with a higher population and higher Internet connectivity are often observed as a top source of web application attack traffic.

ASN and BGP routing as source country indicators / One piece of information that can be used to track attack sources are the Autonomous System Numbers (ASNs), which are assigned to Internet 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 disguise.

In Q1, ASNs also show the US as the top source of malicious web traffic recorded within the Akamai Kona Site Defender infrastructure, followed by Brazil and the Netherlands.

The top three originating ASNs were associated with a VPS farm owned by a well-known cloud IaaS provider. While it is easy to set up a system or VPS in the cloud, it requires technical knowledge to properly to secure it. One misconfiguration or forgotten patch can leave a cloud-hosted system vulnerable. As a result, many systems set up each day can be compromised easily for use in a botnet or other attack platform.

It is also easier, cheaper, and less traceable to set up servers for malicious purposes in the cloud than it is on compromised 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.

Target countries / Once again, the US had the distinction of being both the top source of web application attacks and the top target.

Given that many companies have their headquarters and IT infrastructure in the US, this makes sense. Sixty percent of web application attacks targeted the US, while only 9% targeted Brazil, 6% targeted the UK, and 5% targeted India, as shown in Figure 3-5.

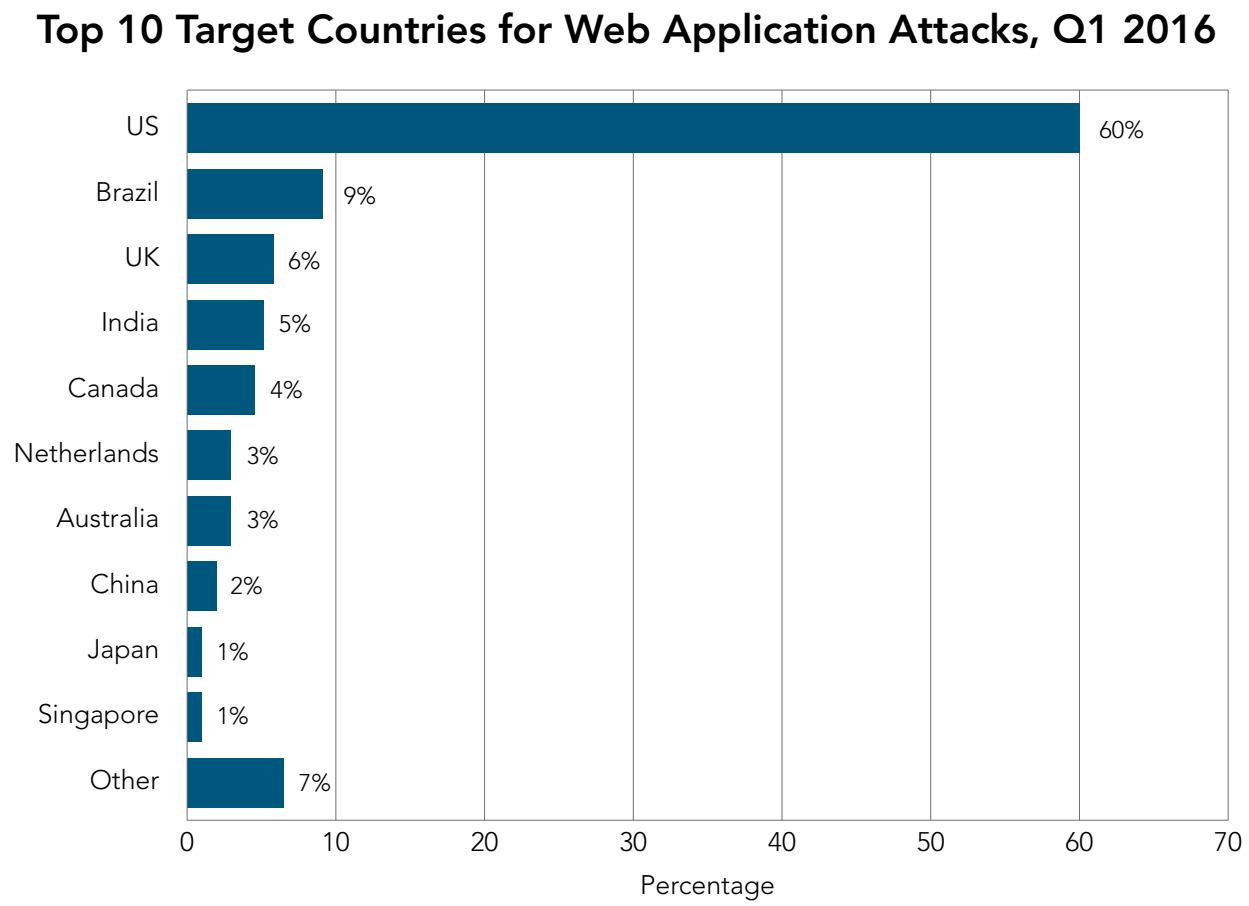

Figure 3-5: US-hosted web sites were targeted six times more often than the second most popular target country, Brazil

3.3 / Web Application Attacks by Industry / This quarter, the retail sector suffered the vast majority of web application attacks: 43% as shown in Figure 3-6. The hotel & travel industry was targeted with 13% of attacks, followed by financial (12%), high technology (9%), media & entertainment (7%), the public sector (3%), SaaS (3%), and business services (2%).

Retail / Retailers are frequently 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. If an adversary is able to find a SQLi vulnerability, the attacker can access the retailer's information. Retailers also have a large number of visitors to their websites. Attackers will find and exploit xss vulnerabilities to deface retailers' websites, causing a loss of trust among customers.

Alternately, the attacker may use a compromised site for a watering hole attack, harming site visitors by, for example, loading malware on their computers. Retailers may also be a target for unvalidated requests. For example, if an attacker could control the price of the item being purchased, items could be sold at a price that is different from what the retailer intended.

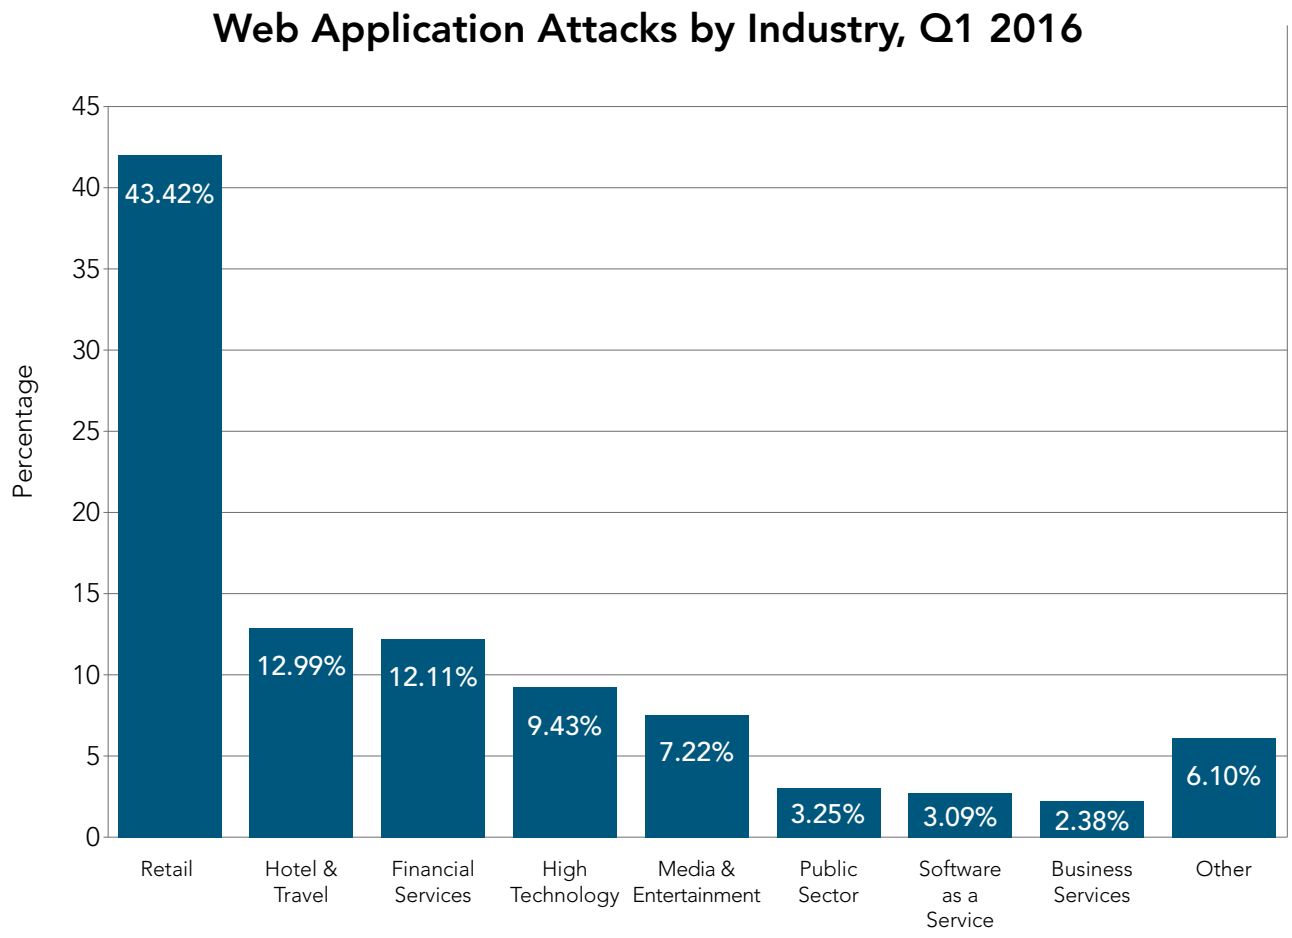

Figure 3-6: As in previous quarters, the retail industry was most frequently targeted with web application attacks in Q1 2016

Media & entertainment / The media & entertainment industry saw a slight drop in attacks: 7% in Q1 compared with 10% in Q3 and Q4 2015. Organizations such as movie studios and news agencies remain attractive targets because they are highly visible and successful attacks on these targets typically generate publicity.

Hotel & travel / The hotel & travel industry saw a slight rise in attacks in Q1: 13%, compared with 10% in Q4 and Q3 2015. This vertical includes hotels, booking agencies, travel sites, and rental agencies. Because many of these organizations are heavily reliant on their online presence to conduct business, any downtime has a major effect. As with retail organizations, travel sites change frequently and have significant amounts of sensitive information. The rate of change means that more opportunities to discover vulnerabilities exist than on more stable sites.

Financial services / The financial industry experienced an increase in Q1 (12%), up from 7% in Q4. Banks and other financial organizations make tempting targets. Even if attackers aren't able to steal money directly, they know they can make a profit through extorting these organizations with the threat of downtime.

High technology / In Q1 2016, the high technology sector suffered 9% of web application attacks — a five-point increase over Q4. This is a broad category encompassing anything from online personnel services to Internet startups.

Figure 3-7 lists the number of attack triggers observed for all classified industries, followed by their percentage of attacks as a whole. The industries not included in Figure 3-6 are shown in red.

This level of granularity is important to understand future attack trends. For example, though the energy and utilities industry only accounted for .06% of web application attack triggers in Q1 2016, the fact that there were 330,757 attack triggers provides a valuable dataset for in-depth research within that industry.

While these other industries do not top the list of targets, they still face substantial and unique risks. By examining them more closely, we can see the beginnings of threats to come.

Web Application Attack Triggers by Industry, Q1 2016

Industry	Attack Triggers	Percentage
Retail	241,821,290	43.42%
Hotel & Travel	72,343,923	12.99%
Financial Services	67,439,442	12.11%
High Technology	52,539,136	9.43%
Media & Entertainment	40,191,100	7.22%
Public Sector	18,116,841	3.25%
Software-as-a-Service	17,227,529	3.09%
Business Services	13,240,770	2.38%
Consumer Goods	7,928,547	1.42%
Gaming	7,919,981	1.42%
Automotive	6,785,434	1.22%
Manufacturing	6,272,125	1.13%
Not-for-Profit	2,509,691	0.45%
Pharma/Health Care	986,849	0.18%
Consumer Services	827,813	0.15%
Energy & Utilities	330,757	0.06%
Real Estate	204,980	0.04%
Education	169,771	0.03%
Miscellaneous	31,276	0.01%

Figure 3-7: 94% of the attack triggers for web application attacks in Q1 2016 targeted just eight industries (shown in black)

3.4 / A Normalized View of Web Application Attacks by Industry /

Akamai has long tracked DDoS attacks at both the application and network layer. DDoS attack statistics are typically the most commented on, reprinted, and discussed numbers that we publish.

Over the years, customers have asked for a similar view into the stealthy application layer attacks that plague enterprises, governments and others; the attacks that organizations such as OWASP have typically tracked and ranked according to prevalence and danger.

But 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 up to 15 – 30% of the world's web traffic at any particular time, the challenge providing this sort of insight to our customers has been threefold: how to store the data we see, how to query it, and finally, how to meaningfully report on it.

Methodology / In the past two years, we've made great progress in tackling the first two challenges. Querying the data has taken a bit more finesse. During the past two years, we've hired a number of data scientists, analysts, and researchers. Today, those researchers make up the Akamai Threat Research team, a team that has set up dozens of heuristics that automatically query the stored data on an hourly basis.

The insight they extract from the data feeds improvements to our Kona Site Defender application protections and our Client Reputation product. The final challenge is reporting on the data. Our reporting methodology divided all Akamai customers into 18 verticals. (Note: For historical reasons, the verticals tracked for application layer attacks are slightly different from those tracked for network layer attacks.)

For each of the customers in these 18 industries, we tracked the number of malicious requests across nine categories of attacks during a 12-week period. The frequency of these attack vectors and the accuracy of the signatures detecting each of the categories were both given weight in the selection of categories.

In order to normalize samples, we removed spikes that we considered to be anomalies in the data. To do so, we calculated the contribution of each sample (individual customer) as a percentage of the total number of attacks per each unique attack

vector. Then we added up the individual percentages for each sample and we finally removed the samples that accrued more than 10% combined, since the data showed those samples (individual customers) were the target of a specific attack campaign.

In the process of normalizing the samples, we discovered some interesting points:

- One particular customer suffered a comprehensive attack campaign and accrued 4% or more of the attacks in seven out of the nine vectors
- Another customer suffered 12% of the total number of attacks observed in the whole quarter
- Two customers together accrued 60% of the Shellshock attacks

After adding up all attacks per vertical and type, we divided the number of attacks in each vertical by the number of customers in that vertical. This eliminates any data skews based on the distribution of our customer base among the verticals.

In Q1 2016, the industries subjected to the largest number of malicious SQLi and LFI requests were the automotive and gaming industries.

SQLi and LFI attacks were attempted against Akamai customers more than any other web application attack vector that we track. The second-most-common web attack vector over HTTP was LFI. These two types of attacks require a very noisy reconnaissance approach. Tools for finding SQLi vulnerabilities can easily make thousands of requests against a site, testing and probing for an entry point. Blind SQLi, which often amounts to asking a site a series of yes or no questions, can require even more requests.

We also observed a prevalence of web application scanners. These point-and-shoot tools are easy to obtain and easy to use. They make a large number of requests when looking for SQLi and LFI vulnerabilities.

The attack triggers by vector and industry are shown in Figures 3-8, 3-9, and 3-10.

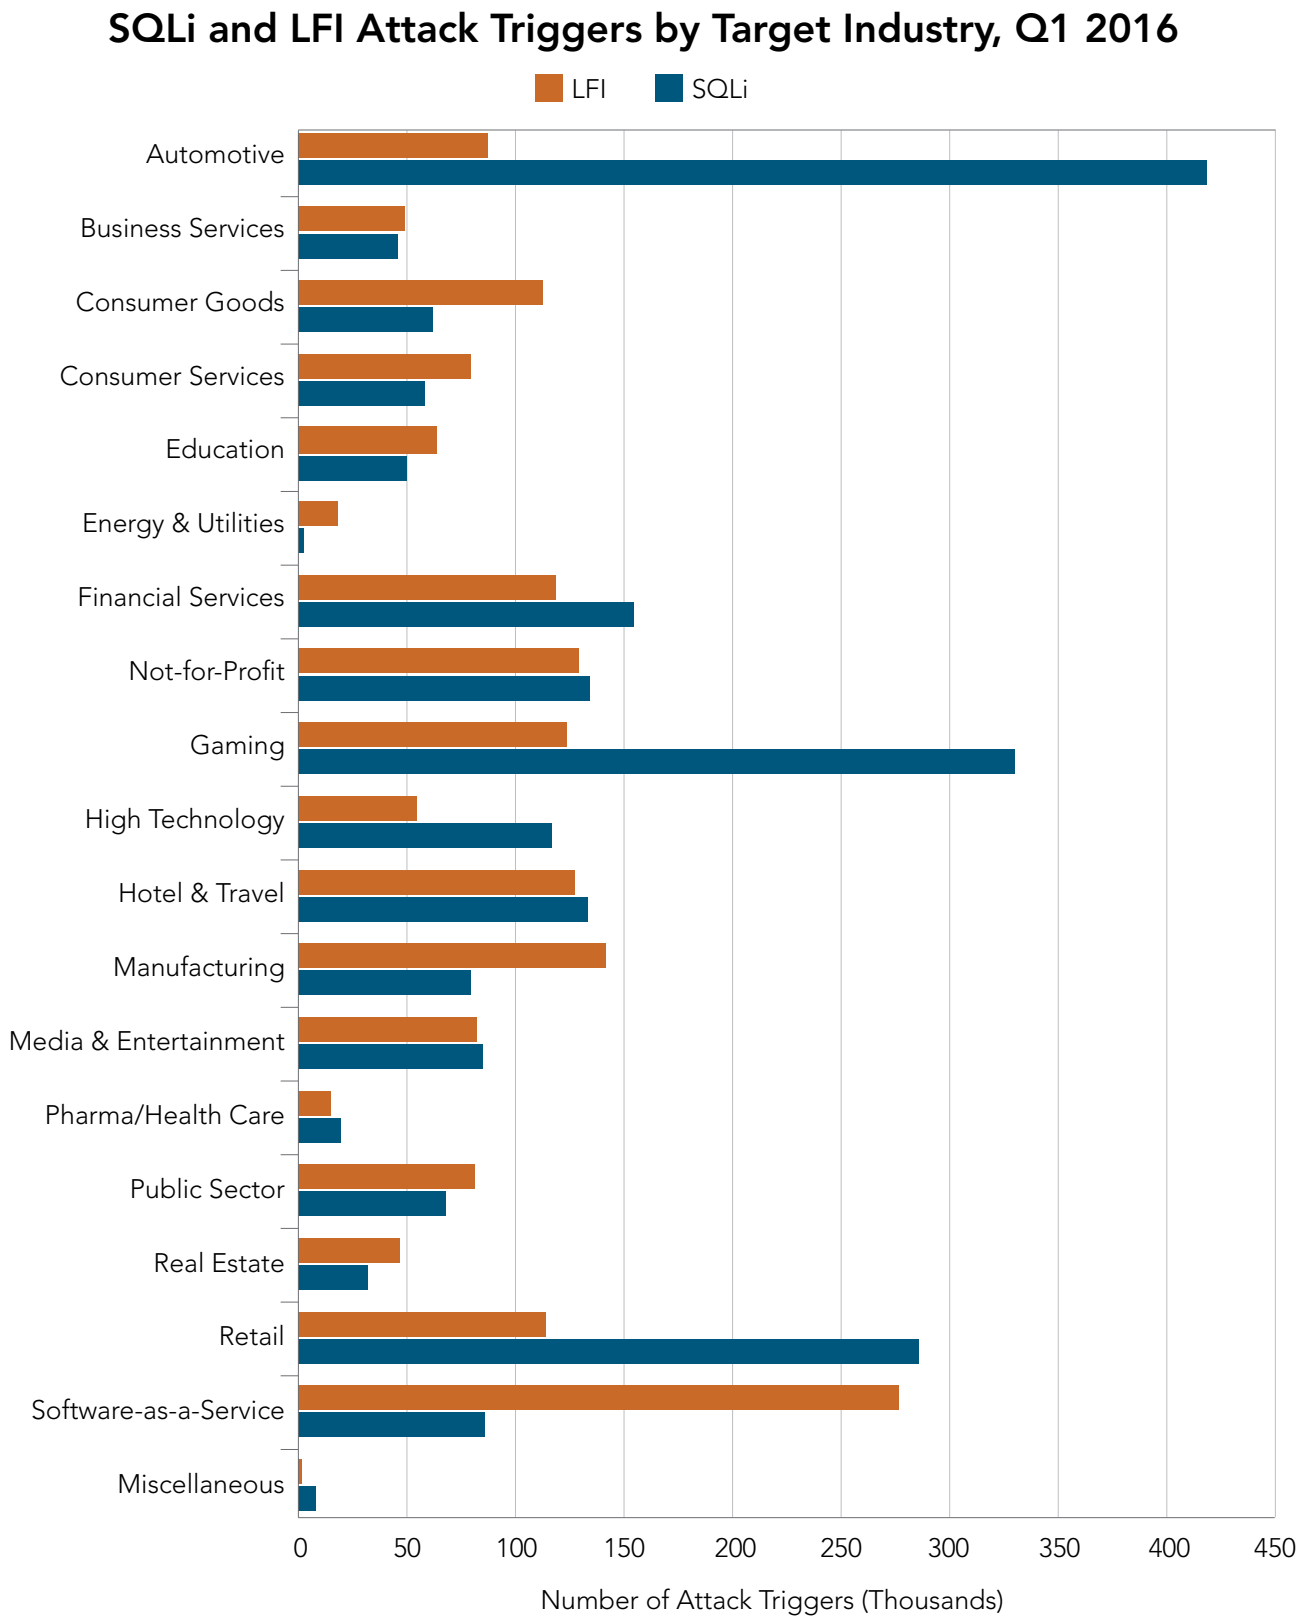

Figure 3-8: LFI attacks were most frequently deployed against the SaaS and manufacturing industries, based on normalized attack data, while SQLi attacks were more frequently observed in the automotive, gaming, and retail sectors

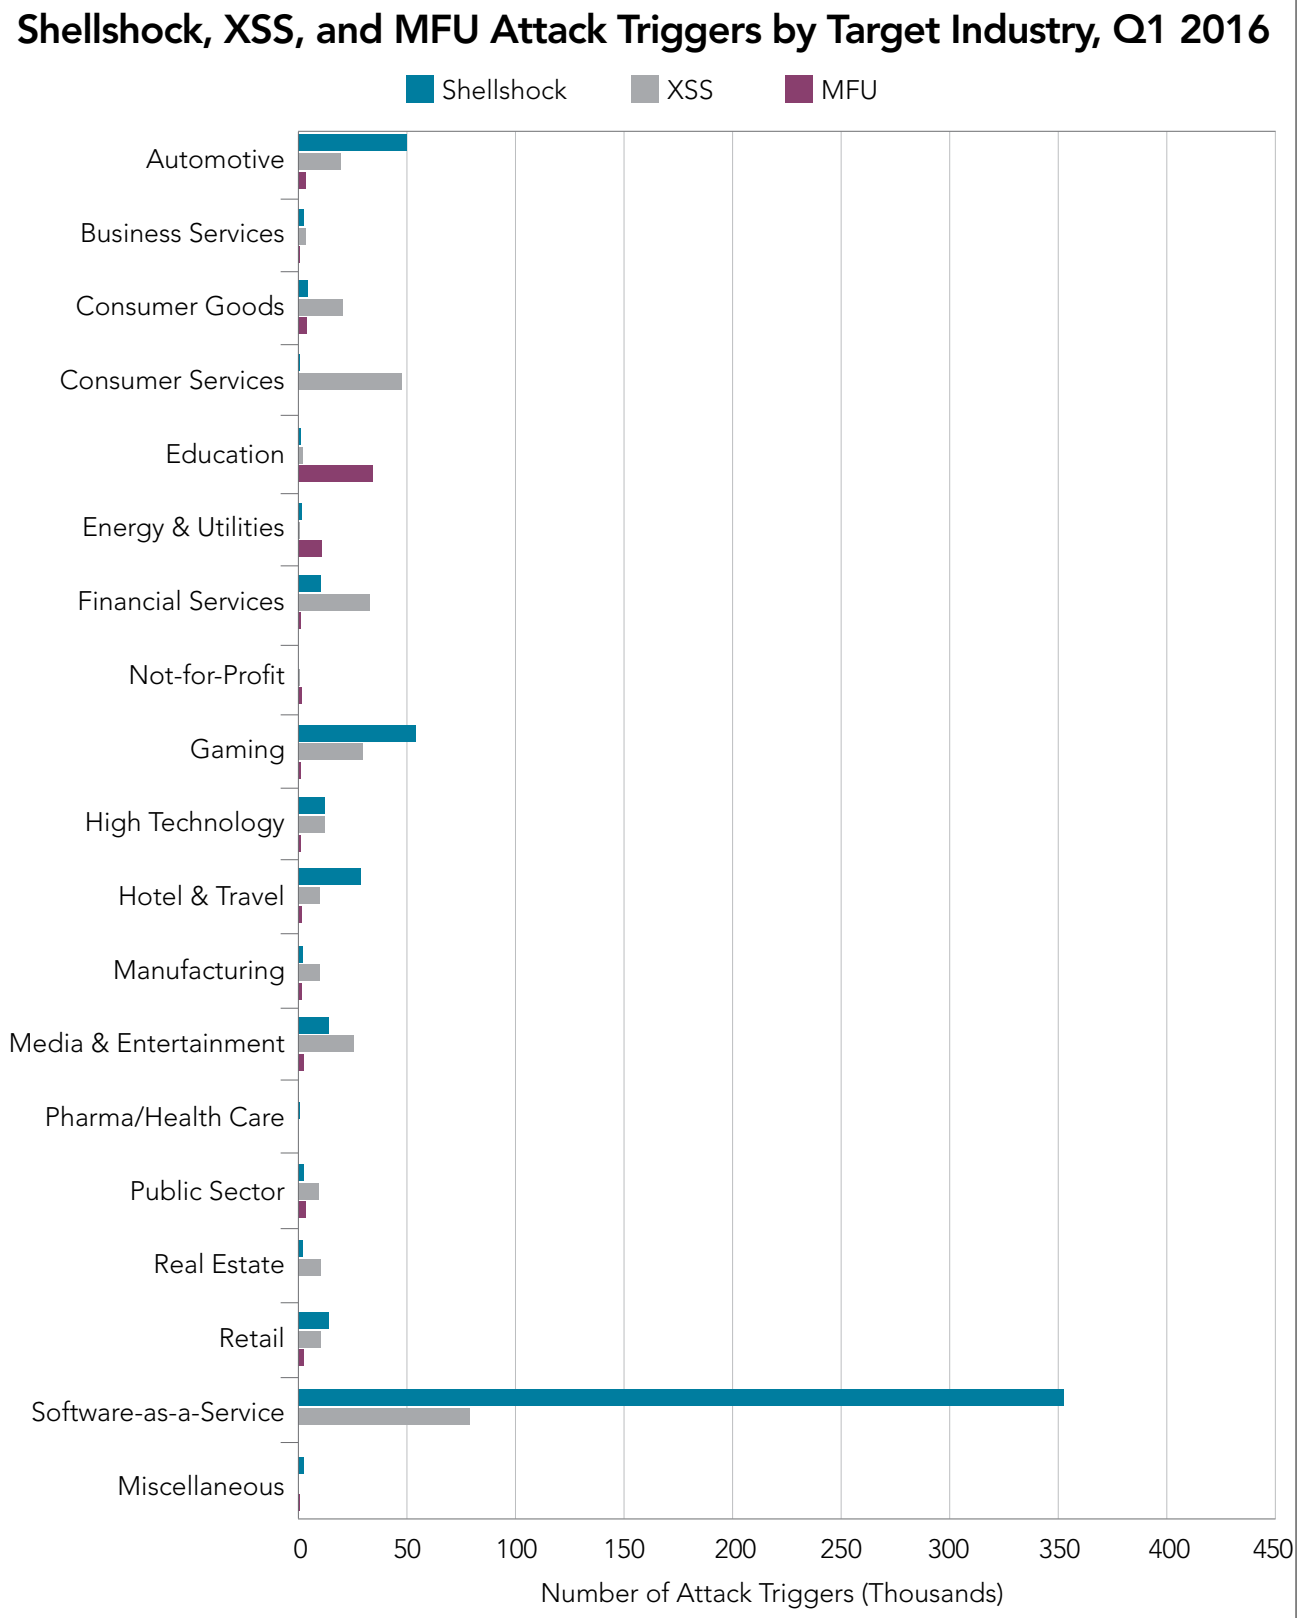

Figure 3-9: The SaaS industry was most frequently targeted with the Shellshock and XSS vectors, based on normalized attack data, while the retail sector was most targeted by MFU attacks in Q1 2016

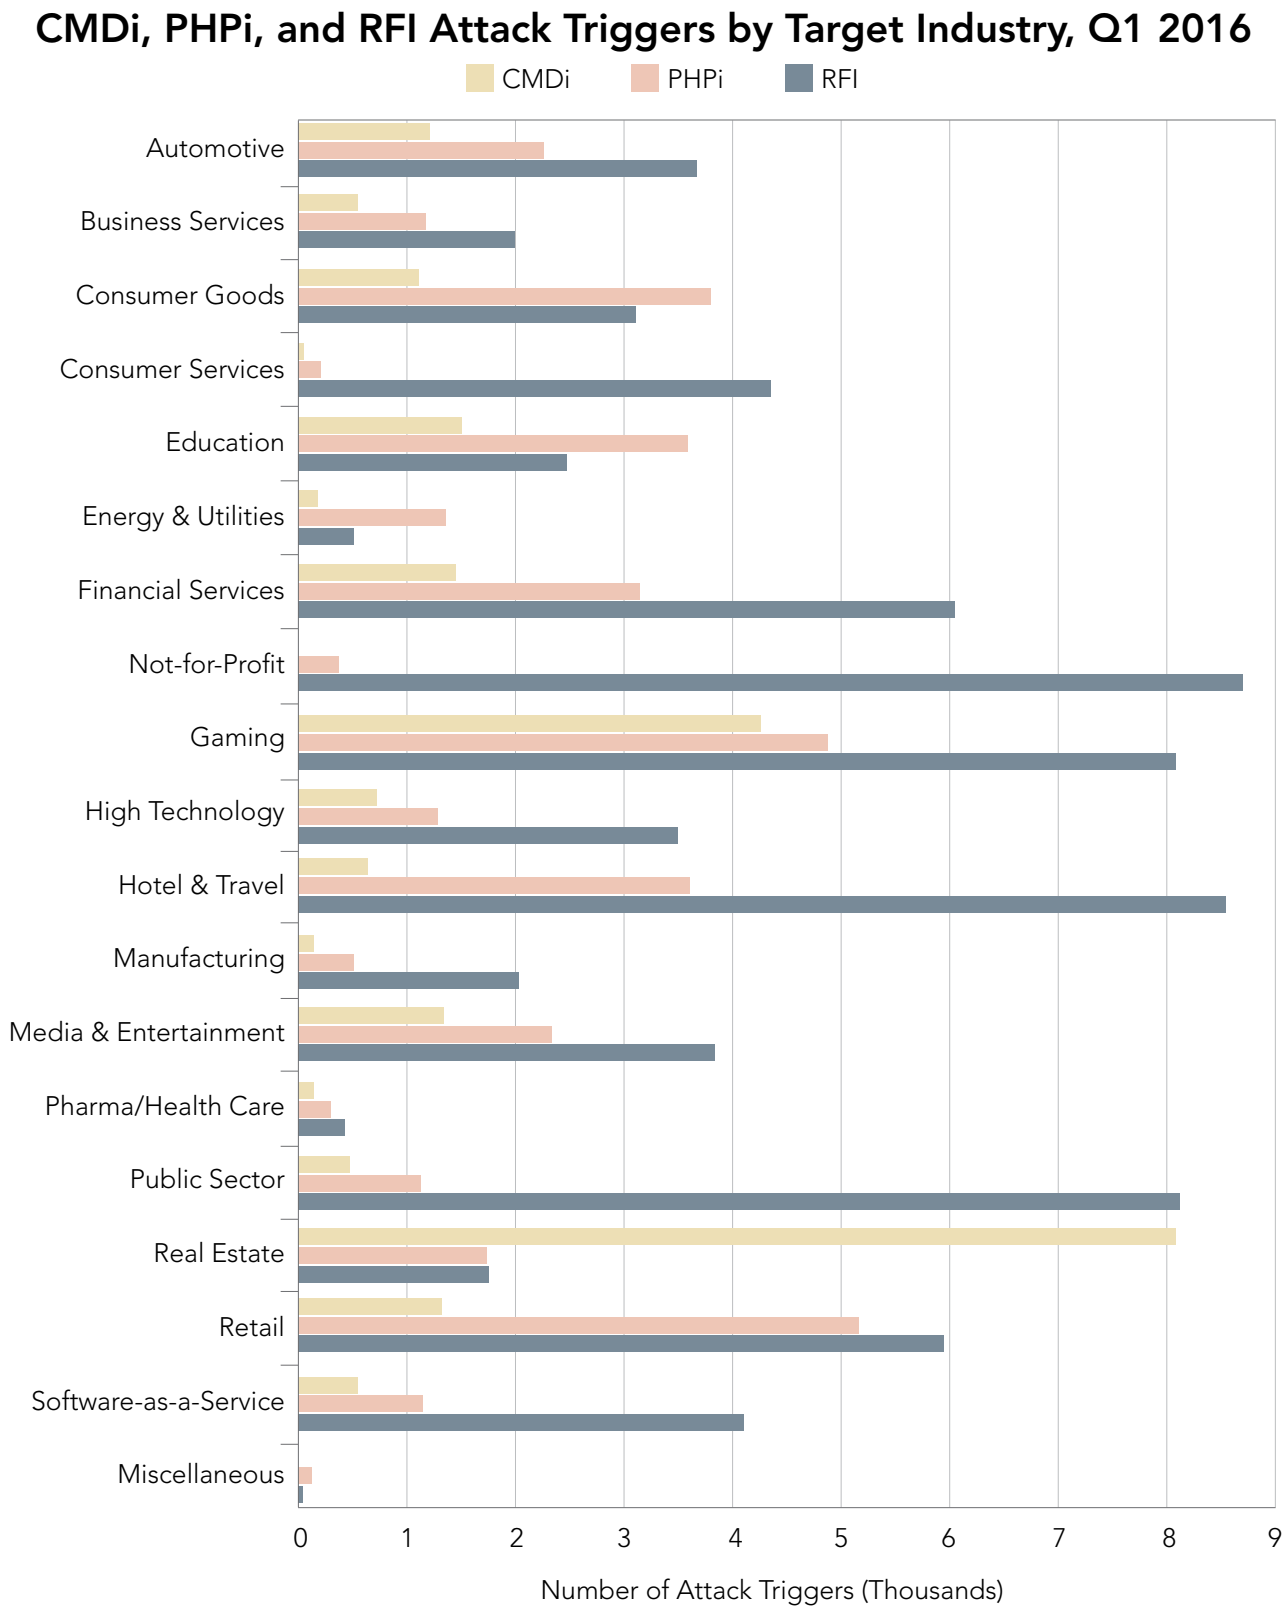

Figure 3-10: CMDi attacks mostly targeted the retail sector, RFI attacks most frequently targeted the hotel & travel and not-for-profit industries, while PHPi attacks mostly targeted the retail and gaming industries, based on normalized attack data

During one specific SQLi attack against an online gaming company's web storefront, the attack appeared to be probing for blind SQLi vulnerabilities by timing query response times. The attack traffic came in small bursts in an attempt to stay under the radar, as shown in Figure 3-11.

```
Connection: keep-alive
    Content-Length: 34
    Accept: */*
    Origin: chrome-extension://cmeakgjggjdlcpncigglobpjbkabhmjl
    User-Agent: Mozilla/5.0 (Windows NT 6.1; WOW64) AppleWebKit/537.36
(KHTML, like Gecko) Chrome/49.0.2623.112 Safari/537.36
    Content-Type: application/x-www-form-urlencoded; charset=UTF-8
    Accept-Encoding: gzip, deflate
    Accept-Language: ru-RU,ru;q=0.8,en-US;q=0.6,en;q=0.4
```

Figure 3-11: In this example, the attacker is probing for blind SQLi vulnerabilities

Throughout the requests collected within this campaign, the attacker mixed cases in an attempt to bypass Web Application Firewall (WAF) rules. This is believed to be the work of a popular web vulnerability scanner, such as Acunetix or Sqlmap. We are also seeing active Account Takeover attacks (ATOs) against major financial institutions.

ATO attacks seem to spike in activity after a major data breach, due to the common practice of password reuse. If a database of 1 million user credentials is stolen, then malicious actors try to automate login attempts against popular sites (such as social media) to see if any credentials are also valid there. In one particular case, the IP address was targeting a main login page and tripped rate controls at 240 messages per second. The machine appeared to be a compromised host machine.

During an LFI scan against a government site, traffic appeared to be generated by the Nikto scanner as shown in the user-agent string in Figure 3-12. The scan originated from three IP addresses that pointed to prominent US-based hosting providers.

Sample Request:

```
Host: target domain.gov
User-Agent: Mozilla/5.00 (Nikto/2.1.6) (Evasions:None) (Test:003023)
Connection: Keep-Alive
Server: AkamaiGHost
Mime-Version: 1.0
Content-Type: text/html
Content-Length: 382
Expires: Thu, 21 Apr 2016 18:15:19 GMT
Date: Thu, 21 Apr 2016 18:15:19 GMT
Connection: close
```

Figure 3-12: This sample request was used to scan for LFI vulnerabilities

3.5 / Bot Traffic Analysis / This quarter we added an analysis of bot traffic collected across the Akamai Intelligent Platform. Figure 3-13 provides a snapshot of more than 2 trillion bot requests observed in one 24-hour period. It excludes bots that are not engaged in DDoS and web application attacks, which are covered elsewhere in this report.

Declared bots represented 40% of bot traffic that day. These bots are operated by legitimate organizations and usually identify themselves by name, homepage, intentions, etc., in the HTTP request (usually within the user-agent header). Those bots serve multiple purposes and services such as search engines, price comparisons, and data analytics.

Detected automation tools and scraping campaigns represented 50% of bot traffic. These bots were detected based on their behavior or some kind of request signature or anomaly. Those bots scrape specific websites or industry segments and don't identify their intentions and origin. In many cases, they impersonate legitimate users or bots.

Bot Traffic, Q1 2016

Overall Bot Traffic

During a full day sample, bot traffic accounted for 30% of all web traffic

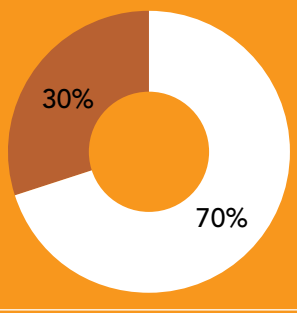

Bot Category Distribution

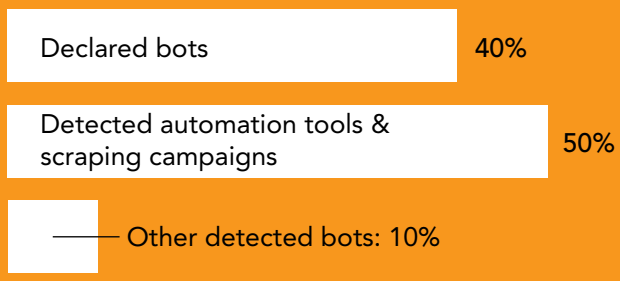

Declared Bots Breakdown / 40% of Bot Traffic

- Web search engines & indexers: 50%
- Media aggregators (social media, news, RSS): 5%
- Commercial aggregators (price comparisons, enterprise data aggregators, scraping enterprise services): 3%
- Analytics & research bots (advertising, SEO analyzers, audience analytics, business intelligence): 15%
- Web monitoring services (performance & health, link checkers): 23%
- Other declared bots: 4%

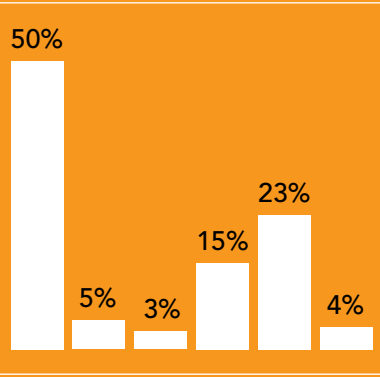

Detected Automation Tools & Scraping Campaigns / 50% of Bot Traffic

- Web-browser impersonators: 55%
- Search-engine impersonators: 8%
- Development frameworks: 30%
- Other detected web scrapers: 7%

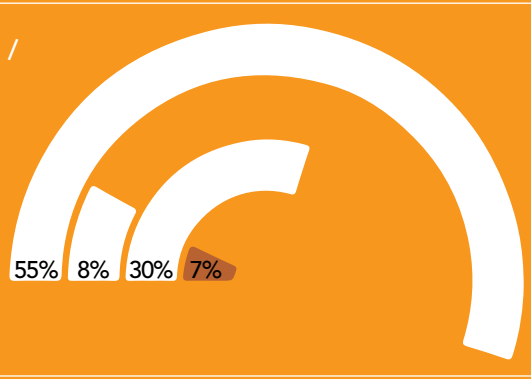

Other Detected Bots / 10%

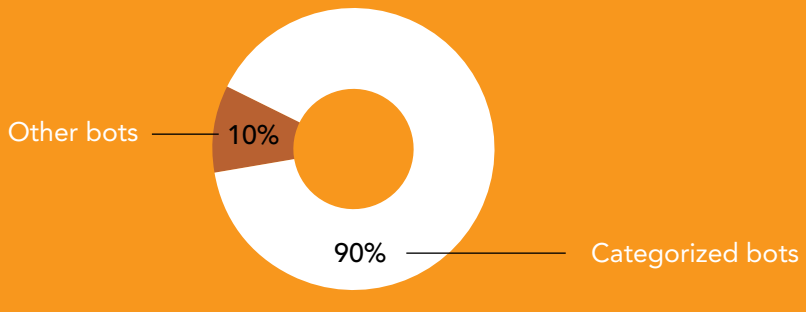

Figure 3-13: A single-day bot traffic snapshot collected on the Akamai Intelligent Platform™, with a breakdown of 24 defined bot categories

OBSERVED BOT TYPES

Web Search Engines & Indexers / These bots are used by search engines to collect and index data. Based on the indexed data, an end user can use search engines to get a ranked set of results based on information need.

Media Aggregators / These bots collect and aggregate data from media resources (TV, music, news, social media, etc.) in order to provide their customers with a centralized and aggregated database based on their needs (e.g., custom homepages and trend analysis).

Commercial Aggregators / These bots collect and aggregate data from commercial and e-commerce sites (shopping, airlines, traveling agencies, etc.). They usually provide their customers with a centralized data source based on their needs (e.g., price comparisons, business intelligence).

Analytic & Research Bots / These bots may target either specific websites/industry segments or a broad range of websites for analysis and research, such as SEO, audience analytics, and advertising.

Web Monitoring Services / These bots approach websites and resources for monitoring, such as link checking, performance testing, and domain name availability.

Other Declared Bots / These bots identify themselves but their origin or intentions can't be defined.

Web-Browser Impersonators / These bots are scrapers that identify themselves as legitimate browsers yet are detected as automated tools.

Search-Engine Impersonators / These bots are scrapers that identify themselves as valid search engine bots in order to access the website.

Development Frameworks / These are scrapers that use HTTP libraries of known development frameworks. They are detected based on request signatures that specify the library name or development framework.

Other Detected Web Scrapers / These bots are detected by scraping behavior or behavior combined with request anomalies.

Other Detected Bots / These bots are detected by some anomalies but their behavior and intention can't be determined.

3.6 / Web Application Security Spotlight: Account Takeover Campaign /

ATO attacks (also known as credential stuffing) use previously breached username and password pairs to automate login attempts. This data may have been previously released on public dumpsites such as Pastebin or directly obtained by attackers through web application attacks such as SQLi. The goal of the attacks is to identify valid login credential data. These credentials may be used by the attackers or be sold to others wishing to gain fraudulent access to user accounts. ATO may be considered a subset of brute force attacks. However, it is an increasing threat because it is harder to identify such attacks through traditional individual account authentication errors. The Akamai Threat Research Team analyzed web login transactions for one week across our customer base to identify ATO attack campaigns.

In ATO attacks, malicious actors typically try to obtain credentials via a direct SQLi attack, a website breach, or even a password dumpsite, as shown in Figure 3-14.

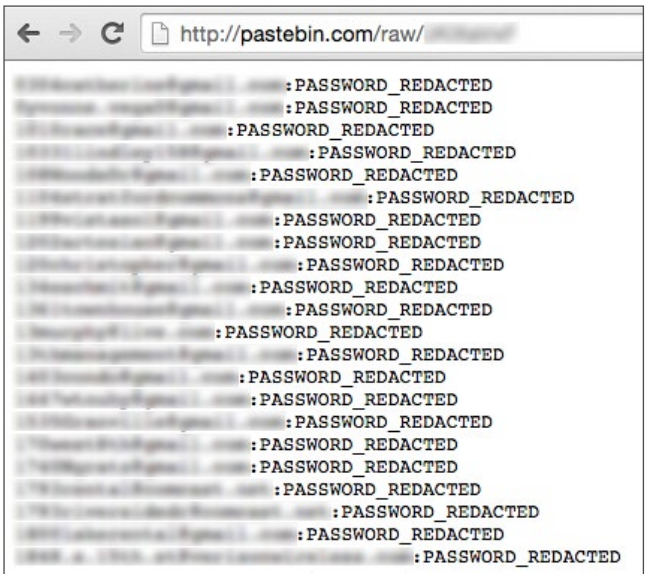

Figure 3-14: Stolen credentials are often dumped on underground sites such as PasteBin

The attacker typically uses an account checker to test the stolen credentials against popular websites such as social media sites or online marketplaces. Most of these account-checking tools have proxy capabilities, which distribute the load across many different source IP addresses. This approach makes rate limiting and blacklisting defenses less effective. An example of an account-checking tool is shown in Figure 3-15.

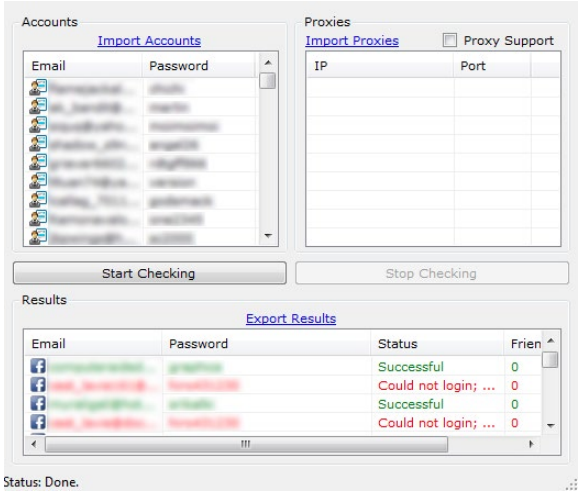

Figure 3-15: Account checking tools are often used to test stolen credentials on social media sites

Once the attacker has validated the credentials, they are typically offered for sale on underground forums and markets, as shown in Figure 3-16. Or, depending on the account and what value it has, the attacker may attempt to cash out value from reward programs and gift cards.

TOTAL ACCOUNTS: (139)

ACCOUNT	COUNTRY	INFO	LOGIN	PASS	PRICE	BUY
	USA	Have Orders	na****	******	2.00	Buy
	USA	Have Orders	ha****	******	2.00	Buy
	USA	Have Orders	ce****	******	2.00	Buy
	WordWide	Fresh Hacked Dell.com Account	do****	******	2.00	Buy
	WALES, HOOK	Fresh Cracked RDP , 2003 Server	sc****	******	6.00	Buy
	WordWide	Fresh Hacked Dell.com Account	ke****	******	2.00	Buy
	USA	Have Orders	ti****	******	2.00	Buy
	WordWide	Fresh Hacked / Upload Checked	NU****	******	3.00	Buy
	WordWide	Fresh Hacked Dell.com Account	wt****	******	2.00	Buy
	FRANCE, BOURGOGNE, DIJON	Fresh Cracked RDP , 2003 Server	re****	******	6.00	Buy
	HARGHITA, ODORHEIU SECUIESC	Fresh Cracked RDP , 2003 Server	sc****	******	6.00	Buy
	WordWide	Fresh Hacked Have Orders	pj****	******	2.00	Buy
	AOL	100k Only Aol , Special Leads	NU****	******	5.00	Buy
	WordWide	Fresh Hacked Dell.com Account	al****	******	2.00	Buy
	WordWide	Fresh Hacked Dell.com Account	ji****	******	2.00	Buy
	USA	Have Orders	rp****	******	2.00	Buy
	WordWide	Fresh Hacked Have Orders	ok****	******	2.00	Buy
	USA	Have Orders	tc****	******	2.00	Buy
	WordWide	Fresh Hacked Dell.com Account	au****	******	2.00	Buy
	USA	Have Orders	me****	******	2.00	Buy
	USA	Have Orders	ka****	******	2.00	Buy
	WordWide	Fresh Hacked Dell.com Account	kt****	******	2.00	Buy
	CALIFORNIA, SOUTH PASADENA	Fresh Cracked RDP , 2003 Server	an****	******	6.00	Buy
	USA	Have Orders	wa****	******	2.00	Buy
	USA	Have Orders	ni****	******	2.00	Buy

Figure 3-16: Once the stolen credentials have been validated, they are offered for sale on underground markets

ATO attack statistics / We analyzed two ATO attack campaigns that took place February 10 – 17, 2016. During this time, many domains were attacked; however, 93% of the attacking IPs were part of a campaign that targeted two specific customers and three domains. These two targeted campaigns were many orders of magnitude greater than all of the other ATO attacks combined.

In the repeated attacks against a customer in the financial services industry, 999,980 IPs were involved in the attacks against the customer's login page. One campaign was responsible for more than 90% of the total attack volume. Here is a closer look at this campaign:

- 993,547 distinct IPs
- 427,444,261 accounts checked
- 22,555 IPs previously blocked based on WAF event logs

The rate of the attack was steady, as 75% of the attackers participated for multiple days, as shown in Figure 3-17.

Attacking IP Persistence — Finance

Number of Active Days	Number of IPs	% of All IPs
1	248,387	25%
2	99,355	10%
3	49,677	5%
4	29,806	3%
5	29,806	3%
6	9,935	1%
7	526,580	53%
Total	993,547	100%

Figure 3-17: Nearly 1 million IP addresses participated in the attack campaign against the financial company*

**Further analysis identified that many of the single day attackers were in fact multi-day attackers, however they either attacked right before or after the one-week time period we initially analyzed. This data is accurate for the stated time period only.*

Entertainment vertical customer / During this same timeframe, 1,127,818 different IPs were involved in attacks. These IPs performed 744,361,093 login attempts and checked 220,758,340 distinct email addresses. The attacks were evenly distributed between two customer sub-domains.

One campaign was responsible for approximately 50% of the total login attempt attack volume.

This campaign lasted all week long, all day long. Here is a closer look:

- 817,390 distinct IPs
- 388,674,528 login attempts
- 65,556,491 email addresses checked

The attacking IP persistence is shown in Figure 3-18.

Attacking IP Persistence — Media & Entertainment

Number of Active Days	Number of IPs	% of All IPs
1	163,478	20%
2	57,217	7%
3	49,044	6%
4	40,869	5%
5	32,696	4%
6	32,695	4%
7	441,391	54%
Total	817,390	100%

Figure 3-18: More than 800,000 IP addresses participated in the attack campaign against the entertainment company

Finance / entertainment campaign overlap / When cross-referencing the attacking sources from both of these targeted campaigns, we identified that 778,786 IPs (more than 70% of the campaign participants) were attacking both customer sites. This implies that our finance and entertainment customers were both targets of a monstrous credential abuse campaign by the same attacking entity, employing the same botnet.

Figure 3-19: Arris cable modems, for home use, were compromised and used in the attack campaigns

Botnet analysis / While many of the participants were hidden behind proxy servers, we did identify an interesting new element: compromised home routers. We clustered the attack IP addresses by geolocation. We noticed a cluster of attacking IPs based out of Mexico and identified that many of the abused systems appeared to be ***Arris*** cable modems, as shown in Figure 3-19. The Arris modem

product line has some ***known backdoors***, where it uses a rotating password of the day and the algorithm has been ***publicly leaked***. We also identified other networking products that were compromised and participating in these ATO attacks, including ***ZyXel routers/modems, as shown in Figure 3-20***.

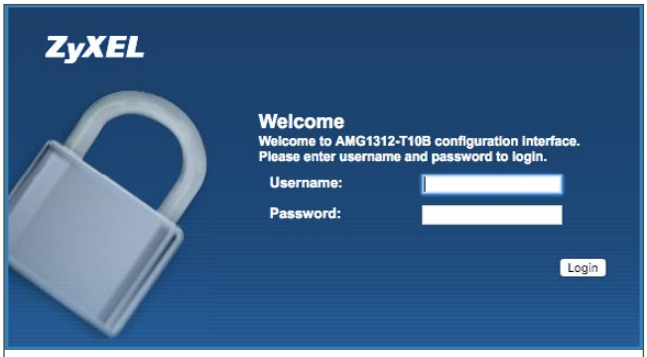

Figure 3-20: ZyXEL home routers were compromised and used in the attack campaigns

As this analysis demonstrates, ATO campaigns are massively distributed and extremely persistent. Organizations need to have an active defense for identifying and mitigating ATO campaigns. The analysis techniques used to identify this ATO activity are actively running within the heuristics applied to the Akamai Cloud Security Intelligence (CSI) platform. This continual analysis of CSI data feeds the Client Reputation product for Kona Site Defender (KSD). In addition to Client Reputation, KSD also offers rate controls that help mitigate brute-force attacks when they exceed thresholds.

[SECTION]4 AKAMAI INTELLIGENT PLATFORM™ FIREWALL ACTIVITY

Akamai firewall data from the platform perimeter provides a broad look at attack activity at the global platform perimeter — with information on attack traffic coming from more than 216,000 servers in more than 122 countries and within more than 1,545 networks around the world. This samples the background radiation of the Internet as well as malicious traffic attacking our services.

At the platform perimeter, 2 PPS per system are logged and analyzed, giving us a more accurate, broader look at affected hosts and attack tactics. This data creates a larger lens to examine the types of non-layer 7 attacks being attempted against Akamai customers.

Scanner activity shared in our Q4 2015 report demonstrated that malicious actors use scanners and probing to perform reconnaissance on their targets before launching attacks. There is also a correlation between new and existing methods of reflection, and the number of scans for suitable reflectors on their well-known ports, as shown in Figure 4-1. Many tools used in these attacks are paired with a scanner script that sends out queries across the Internet, similar to the attack tool query payload. Hosts that respond as expected are logged for inclusion in a list for use as reflection sources in later DDoS attacks.

Reflector Target Services and Port Numbers

Service	Port
QOTD	17
CHARGEN	19
RPC	111
NTP	123
SSDP	1,900
TFTP	69

Figure 4-1: Service port numbers of tracked reflectors

Reflection attacks / For this section, we focused on UDP-reflected DDoS attacks, including SSDP, NTP, CHARGEN, Quote of the Day (QOTD), TFTP, and RPC. Figure 4-1 lists the services and associated port numbers of the reflectors we tracked. Sentinel has been a less leveraged vector and is not expected to outlive some of the more pervasive protocols used as reflectors. This low activity leaves it out of the Q1 2016 dataset. Replacing sentinel is the newly popular TFTP vector.

By looking at the top reflection sources by ASN, we saw that many reflector sources are based in Asia. The top two ASNs located in China made up 40% of the top 10 ASNs by unique IP count, as seen in Figure 4-2. These two ASNs accounted for more than 32% of the SSDP reflectors. The first and second top ASNs included more than 20,000 and 15,000 unique reflector sources respectively.

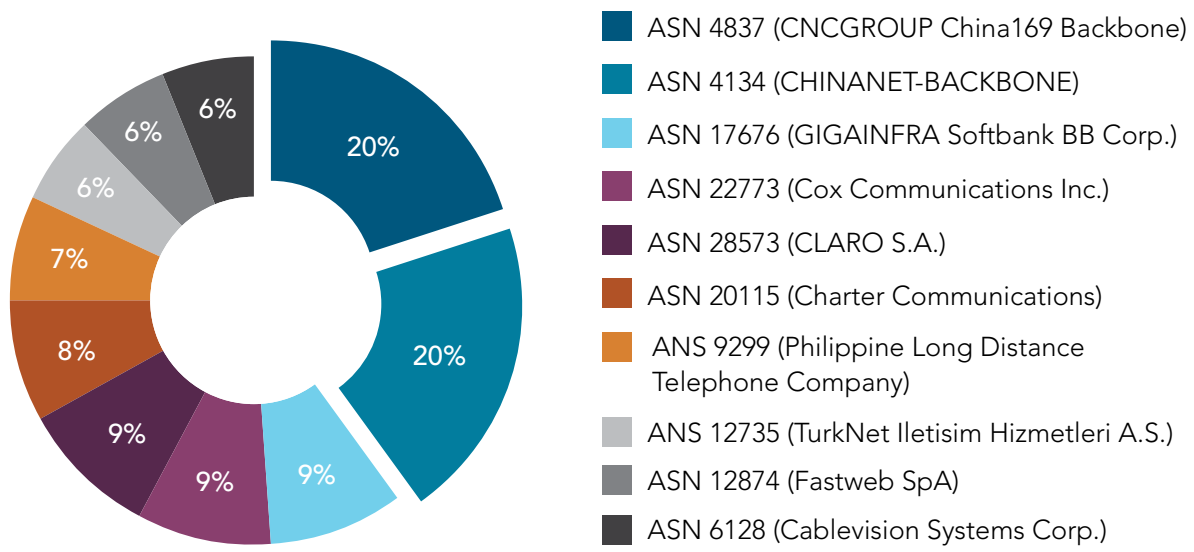

Figure 4-2: The top two ASNs used in reflection attacks, both based in China, accounted for 40% of recorded attack traffic

While most SSDP reflector sources tend to be from home connections, it is interesting to note the diminished use of SSDP in reflected DDoS attacks while NTP continues to rise — it is now second to DNS in terms of reflection DDoS attacks. NTP, CHARGEN, and QOTD reflectors are generally located at cloud hosting providers where those services run.

Cloud-hosting locations would typically support higher-bandwidth attacks. However, the increasing bandwidth available in the home may change that. When looking at SSDP vs. NTP in single-vector attacks, the trend was toward smaller NTP attacks while SSDP slightly trended toward larger attacks.

Many NTP reflectors don't respond correctly to the usual `monlist` query sent in rapid succession during a DDoS attack. In these cases, the attacker is probably burning more bandwidth than the reflector is sending back to the target, reducing the amplification value of their attack. Most booter/stressors operators create lists of NTP and other reflector IPs, but it's not clear how diligently these lists are maintained.

Many attack signatures also reveal that NTP and other attacks are returning ICMP port `unreachable` messages, which indicate the server is no longer running or the NTP port is being actively filtered to prevent further abuse.

Figure 4-3 displays the most prevalent areas for the SSDP, CHARGEN, NTP, and QOTD attack sources identified in Q1 2016. It was populated by logs identifying more than 600,000 reflectors, as shown in Figure 4-4. This was a jump from the 525,850 unique reflectors in Q4 2015. During Q1 2016, most of the hosts and devices leveraged in reflection attacks continue to be clustered around the US, Asia, and Europe.

DDoS Reflector Heat Map, Q1 2016

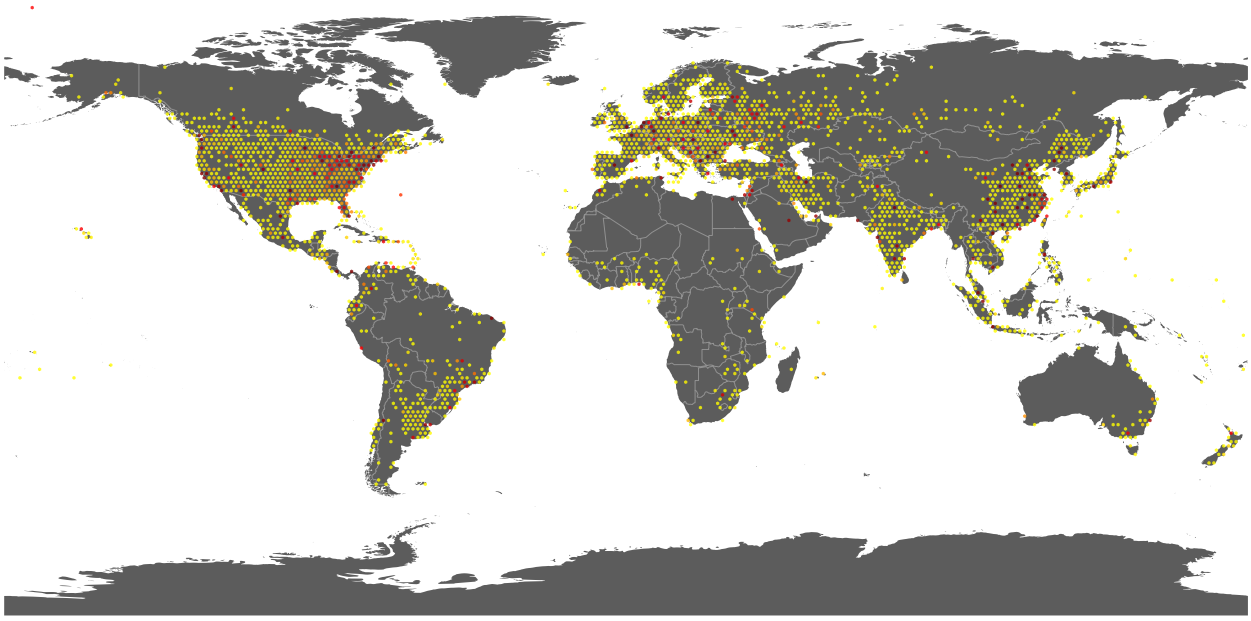

Figure 4-3: The location of leveraged Internet devices used in reflection-based DDoS attacks during Q1 2016 was concentrated in the US, Asia, and Europe

NTP saw a surprising 71% rise in sources reflecting NTP. Again, many seem to be failed attempts to reflect off unwilling NTP servers. Overall, this puts NTP reflection sources at the top of the vector list for the first time, accounting for 59% of all sources logged, as shown in Figure 4-5. While SSDP sources plunged 46% from the previous quarter, sources of NTP, CHARGEN, and QOTD rose by more than 65% apiece, as shown in Figure 4-6.

DDoS Reflectors Unique IP Count by Vector, Q1 2016

Vector	IP Count
NTP	359,443
SSDP	114,270
CHARGEN	47,491
RPC	34,403
QOTD	32,532
TFTP	21,439

Figure 4-4: More than 600,000 DDoS reflectors were tracked across Akamai's Intelligent Platform™ in Q1 2016

DDoS Reflection Sources, Q1 2016

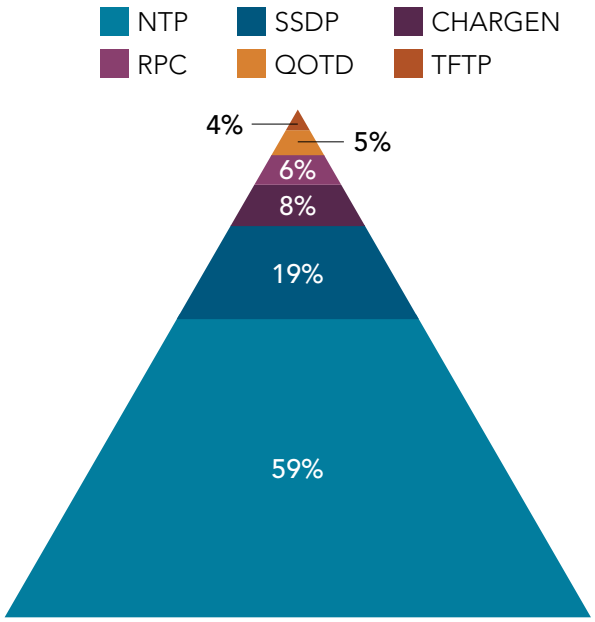

Figure 4-5: NTP was the top source of reflection DDoS attacks (59%), followed by SSDP (19%) in Q1 2016

Changes in Reflector Type, Q1 2016 vs. Q4 2015

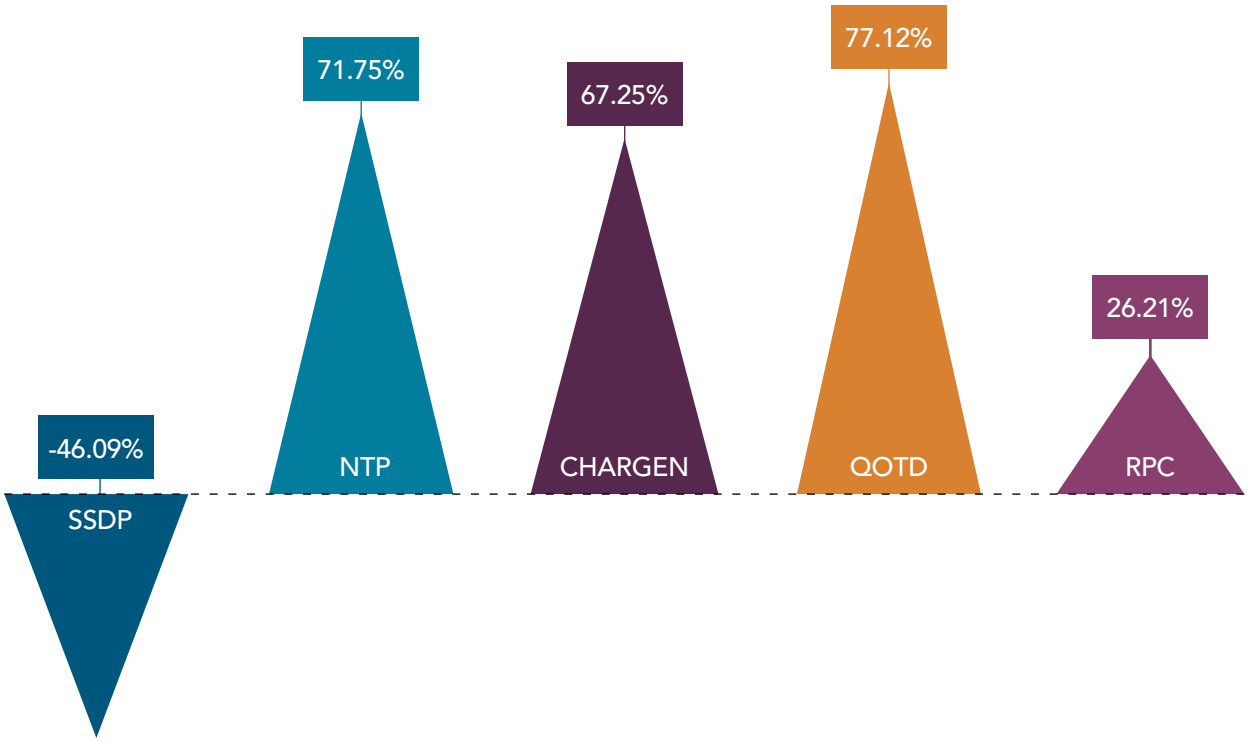

Figure 4-6: The use of SSDP as a source of reflection attacks dropped by almost half from the previous quarter, while the use of NTP, CHARGEN, and QOTD rose by more than 65% apiece

[SECTION]5 CLOUD SECURITY RESOURCES

Akamai released eight threat advisories, vulnerability updates, and attack case studies in Q1 2016.

5.1 / *Scraper and Bot Series — When Good Bots Go Bad* / In March, Akamai launched a report series about bots and scrapers, based on continued research by Akamai's SIRT. In the first installment, we discussed the various types of bots and scrapers we have encountered, and how you could react to each. This paper focused on the known so-called good bots: traffic that is encouraged because it can be helpful to a business.

Bots are programs designed to do specific tasks on the web. Common bot purposes include denial of service, service checking, and file transfer, among others, both malicious and legitimate.

Web scrapers collect web content. A scraper makes requests to web pages and processes the collected data. This data could be hotel or airline prices, store locations, current sales, or any other valuable data. This stored information is analyzed and used for competitive intelligence, or sold to other parties.

IP blocking is rarely an effective defense, but the majority of malicious bots can be stopped by rate limits. Targeted bots, however, require careful research and additional planning to defeat. Akamai recently launched ***Bot Manager*** so customers could manage bot traffic more effectively.

Akamai observes bots of all types every day against virtually every industry. The type of bot is often dependent upon the industry. At a minimum, every site will have search engine bots making requests to their site. Some studies have shown that more than half of all Internet traffic is from automated sources. While some of this is welcomed by site owners, much of it is not.

Our next publication in this series will discuss malicious bots, the type that try to harm sites by using DDoS and related tactics. After that, we will talk about highly aggressive scraping bots, which try to gather information as quickly as possible. Traffic from these bots may superficially appear to be a DDoS attack. It may even actually cause a denial of service. Finally, we will wrap up the series with a paper on stealthy bots that utilize a slow and low methodology to avoid detection. These bots are often account checkers or card validators looking to find valuable information to sell on the black market.

This series is based on years of research by Akamai and explains how to handle each type of adversary.

5.2 / *#OpKillingBay Expands Attacks* / Operation Killing Bay, better known as #OpKillingBay on social media sites, expanded its attacks in Q1. Historically, attackers participating in OpKillingBay have targeted Japanese government websites and sites of companies participating in whale and dolphin hunting. These attacks are led by activists or protesters using hacktivism to spread their message and harm their chosen opponents.

Recently however, Akamai SIRT saw a shift in tactics with OpKillingBay. While they are still attacking the same types of sites as in the past, they are also attacking sites unrelated to the hunts. In fact, one member of OpKillingBay stated that a company was targeted not because it supported the hunts, but because it did nothing to oppose them.

Akamai also observed the group threatening to attack whaling groups from other parts of the world. We have uncovered target lists for sites in Denmark, Iceland, and the Faroe Islands. We also saw that the attackers involved in OpKillingBay have spread out and participated in other operations as well. OpKillingBay has persisted for approximately three years.

On Jan. 12, we started to observe attacks against a Japanese automotive web site. Another automotive company joined the list of victims on Feb. 4, and attacks continued every few days.

At this point, it's unclear if the person claiming credit is actually the one launching the attacks, but Akamai has viewed the attack traffic against the mentioned organizations, and others, that have been targeted within our infrastructure.

The most-targeted industries thus far have been seafood companies, government agencies, and theme parks. The group has declared any government site from Japan or Iceland to be a target, given those countries' whaling histories.

5.3 / *BillGates Malware Used in DDoS Attacks* / We continued to see the BillGates Trojan/bot family of malware being used to launch DDoS attacks in Q1. Attackers who control the malware, as first disclosed on a Russian IT website in February 2014, can gain full control of infected systems.

The attack vectors available within the toolkit include ICMP flood, TCP flood, UDP flood, SYN flood, HTTP flood (layer 7), and DNS reflection floods. This malware is a descendant of ***Elknot's malware*** source code. It has been detected in the wild for a few years and the botnets that compose it have grown and are launching significantly larger attacks.

There is a possibility that after the takedown of Xor botnet, the malware actors began using different means and/or different botnets to continue their onslaught of attacks directed at the same primary group of targets. This activity has been observed by Akamai over the last six months.

The botnet targets are similar to previously confirmed Xor botnet attacks, with most of the targets being online gaming institutions in Asia.

Akamai SIRT observed inactivity from an Xor C2 in Q4 2015, which was publicly announced by a third-party source and believed to be part of a takedown operation. Once that occurred, we believe the attackers started using the BillGates botnet to launch attacks against the same target list.

The ***advisory*** included validated DDoS attack campaigns Akamai mitigated and an example of its use as one of the source tools in combination with the use of a booter site. The advisory also described the detection of malware infections, identified attack patterns, and explained how to clean an infected machine.

5.4 / *Akamai Responds to Forwarding-Loop Issue* / In late February, Akamai became aware of the research paper titled, "***Forwarding-Loop Attacks in Content Delivery Networks (CDNs),***" published by Jianjun Chen et. al on Feb. 29.

We reviewed the researchers' findings, and were confident that we already had adequate counter-measures in place to thwart any attempt to use Akamai as an attack vector in the manner described by the paper.

The paper described four types of forwarding-loop attacks against CDNs: self-loop, intra-CDN loop, inter-CDN loop, and dam flooding. The paper acknowledged that Akamai was not vulnerable to the first two. The third attack (the inter-CDN loop attack) was described as a looping between multiple CDNs. Finally, the fourth, dam flooding, was described as coupling "forwarding-loop attacks with timely controlled HTTP responses to significantly increase damage."

While Akamai does not publically disclose or discuss our security countermeasures, we reiterate that we had and still have sufficient countermeasures in place to detect and defend against all these attacks, as well as substantial capacity to absorb traffic spikes.

5.5 / *IKE/IKEv2 Ripe for DDoS Abuse* / Throughout Q1, Akamai SIRT conducted research into the security posture of the Internet Key Exchange (IKE & IKEv2) protocol. We released a paper in February outlining the findings thus far, including choices in the protocol itself that attackers could potentially leverage to launch reflected DDoS campaigns.

Our motivation to examine it was based on the prevalence of IKE/IKEv2, which is used to facilitate secure key exchanges between peer devices in the IP Security (IPsec) protocol suite. It is widely deployed in multiple secure tunneling applications such as VPN products from major vendors and open source projects.

Over the past four years of active monitoring and threat advisory releases concerning reflection-based DDoS attacks, several abusable UDP protocols have come to our attention. Results from this research have informed *Akamai's State of the Internet / Security* reports about trends in the DDoS threat landscape. This history has

sparked efforts internally to help discover new potential UDP-based reflection and amplification vulnerabilities, with the goal of disclosing, cleaning up, and fixing issues before they can be weaponized for DDoS.

The ***full paper*** delves into the history of IKE, offers visuals to illustrate where the weaknesses are, and offers steps organizations can take to reduce risk exposure.

5.6 / *Akamai and the Glibc Vulnerability (CVE-2015-7547)* / In February, Akamai investigated the Glibc vulnerability outlined in ***CVE-2015-7547*** to see how our systems may have been affected.

As part of the DNS query process, Glibc is used by many systems across the Internet — and at Akamai. All versions of Glibc's `getaddrinfo()` library functions, since version 2.9, are potentially vulnerable to a range of attacks based on a stack buffer overflow.

Researchers from Google and Red Hat discovered the vulnerability, and ***a patch*** was issued.

Akamai quickly patched our internal systems. Most of the company's external systems use a local resolver and are not vulnerable to attacks against `getaddrinfo()` on the wire.

5.7 / *Akamai and the DROWN Vulnerability* / In February and March, Akamai hardened systems against the Decrypting RSA with Obsolete and Weakened eNcryption ***(DROWN) vulnerability*** (***CVE-2016-0800***), which exploits a legacy encryption protocol, SSLv2, in order to compromise keys that might have been reused by secure modern protocols such as TLSv1.2. The modern protocols do not leak the SSL/TLS keys themselves.

We took the necessary steps to protect both our customer-facing and critical internal systems from this vulnerability as of March 1. We started a process to identify and patch non-critical systems on an as-needed basis.

The DROWN attack, ***described here***, allows an adversary to compromise secrets from modern-TLS connections if any machine will accept SSLv2 connections using the same key and certificate.

Our secure delivery services were not vulnerable to DROWN, but individual customers have the option to enable SSLv2 for their own sites. Doing so would expose that customer's connections to DROWN.

While Akamai's secure delivery provides protection, customers were still advised to verify that the origin servers they operate themselves do not use SSLv2.

5.8 / *DNSSEC Targeted in DNS Reflection, Amplification DDoS Attacks* / During the past few quarters, Akamai has observed and successfully mitigated a large number of DNS reflection and amplification DDoS attacks abusing DNSSEC-configured domains.

As with other DNS reflection attacks, malicious actors continue to use open DNS resolvers for their own purposes, effectively using these resolvers as reflectors. This technique has also been linked to the DDoS-for-hire underground market.

The attacks were outlined in a ***security bulletin*** written by Akamai SIRT.

DNSSEC is a suite of Internet Engineering Task Force (IETF) specifications for securing certain information provided by DNS. It is essentially a set of extensions to DNS that provides origin authentication of DNS data, data integrity, and authenticated denial of existence. These additional security controls are designed to protect the Internet against certain types of attacks. There is a ***full list of all requests for comments*** associated with DNSSEC.

To date, Akamai has observed several domain names utilized for these attacks. Although the domains listed in the bulletin have been used in these attacks, the list is not exhaustive.

Since the beginning of November 2015, Akamai has detected and mitigated more than 400 DNS reflection/amplification DDoS attacks using a variety of domain names implementing DNSSEC. DNSSEC prevents attackers from crafting believable false DNS responses, such as might be used to lure a user to an attacker's own website. This extra security offered by DNSSEC comes at a price, as attackers can leverage the larger domain sizes for DNS amplification attacks.

The highlighted domain was observed in DDoS attacks against customers in multiple verticals over the same time period, and based on our investigations, we believe these attacks are most likely the work of attackers making use of a DDoS-for-hire service that uses purchased VPS services, public proxies, a classic botnet, and basic attack types such as DNS reflection attacks, SYN floods, UDP floods, SSDP floods, NTP floods, ICMP floods, and even HTTP GET floods.

The report goes into detail about individual attacks, including screenshots and other graphics, and outlines steps organizations can take to protect themselves.

5.9 / *Akamai Customers Not Vulnerable to SLOTH* / In January, Akamai was informed of a new TLS vulnerability — Security Loss due to the use of Obsolete and Truncated Hash (SLOTH) — by researcher Karthik Bharghaven. Akamai worked with the researcher to confirm and fix the vulnerability in an expedient manner prior to public disclosure. Consequently, we minimized the chances of an exploit and have determined that Akamai customers are now not vulnerable to SLOTH.

SLOTH is a tricky attack against SSL/TLS connections that uses various protocol interactions in unexpected ways. The key point is that weak digests are used in SSLv2 (MD5) and that later versions of the protocol can be tricked into using those digests. It's important to know that the digests are used to protect the SSL/TLS messages exchanged between client and server.

Akamai eliminated the vulnerability to SLOTH quickly to help best protect our customers and our platform.

For an interesting graphic on message digest lifetimes, see ***http://valerieaurora.org/hash.html***. For the full technical details on SLOTH, see ***http://www.mitls.org/pages/attacks/sloth***.

5.10 / *How Web Applications Become SEO Pawns* / In the second half of 2015, Akamai's Threat Research team identified a sophisticated Search Engine Optimization (SEO) campaign that uses SQLi to attack targeted websites. An ***advisory*** on the subject was released in January.

Affected websites distribute hidden Hypertext Markup Language (HTML) links that dupe search-engine bots and skew page rankings to the point where they're no longer accurate.

Over the course of a two-week period in Q3 2015, Threat Research analyzed data gathered from the ***Akamai Intelligent Platform™*** and saw attacks on more than 3,800 websites and 348 unique IP addresses participating in the various campaigns.

Key points:

- When searching the Internet for the HTML links that were used as part of this campaign, Threat Research identified hundreds of web applications containing these malicious links.
- When searching for a combination of common words such as *cheat* and *story*, it was apparent that the cheating stories application appeared on the first page of the leading search engines.
- Threat Research looked at Alexa analytics and the ranking of the cheating stories application dramatically increased during the three-month span.

[SECTION]6
LOOKING FORWARD

In the coming months, we expect to see more records set for the number of DDoS attacks on Akamai's routed network, driven in large part by the continued use of stresser/booter botnets. Though the attack vectors and methods will continue to vary, the majority of attacks will likely rely on reflection. There's little chance of a rapid cleanup of the servers that enable these attacks. The inclusion of new vectors such as TFTP reflection also make cleanup a constantly changing task. As we've seen in recent quarters, the number of targets attacked will likely grow incrementally, while the number of attacks is expected to grow by leaps and bounds, leading to large increases in attacks per target.

Readers can expect us to focus more on analysis of traffic based on the assigned ASN in association with its BGP routing, looking to identify major sources of malicious traffic.

We expect the US and China to remain the top sources of malicious traffic because of the sheer number of devices, vulnerabilities, and users in these countries. But there will be the occasional surprise, such as the UK taking the top spot in Q3 2015 and Turkey in second place last quarter. It is likely that cloud providers will remain the biggest trouble spot unless they do more to improve their default system configuration security procedures.

Distributed reflection Denial of Service (DrDoS) attacks will remain a popular weapon of choice for attackers, though it remains to be seen if vectors like NetBIOS, RPC PORTMAP, and now TFTP servers become as prevalent in reflection DDoS attacks. Surprisingly, despite a decreasing number of ideal available resources, NTP reflection surged near the end of Q4 2015 and continued into Q1.

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

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.

One driver for future threats is the continued proliferation of easy-to-use DDoS-for-hire technology. The same technologies that make the user experience easier for law-abiding people will also create an easier experience for the online criminal community.

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 construction, and monetization of cyberattack frameworks.

***State of the Internet / Security* Team**
David Fernandez, Akamai SIRT
Bill Brenner, Akamai SIRT
Jose Arteaga, Akamai SIRT
Ezra Caltum, Threat Research Unit
Martin McKeay, Sr Security Advocate
Dave Lewis, Security Advocate
Jon Thompson, Custom Analytics
Ryan Barnett, Threat Research Unit
Larry Cashdollar, Akamai SIRT
Miguel Serrano, Security Marketing
Ory Segal, Threat Research Unit
Yossef Daya, Threat Research Unit

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

Contact
SOTIsecurity@akamai.com
Twitter: *@akamai_soti* / *@akamai*
www.akamai.com/StateOfTheInternet

Akamai FASTER FORWARD

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*.

©2016 Akamai Technologies, Inc. All Rights Reserved. Reproduction in whole or in part in any form or medium without express written permission is prohibited. Akamai and the Akamai wave logo are registered trademarks. Other trademarks contained herein are the property of their respective owners. Akamai believes that the information in this publication is accurate as of its publication date; such information is subject to change without notice. Published 06/16.