

[state of the internet] / security

Q2 2017 Report

AT A GLANCE

Web application attacks, Q2 2017 vs. Q2 2016

25% increase in total web application attacks
86% increase in attacks sourcing from the U.S. (current top source country)
60% decrease in attacks sourcing from Brazil (Q2 2016 top source country)
44% increase in SQLi attacks

Web application attacks, Q2 2017 vs. Q1 2017

5% increase in total web application attacks
4% increase in attacks sourcing from the U.S.
21% increase in SQLi attacks

DDoS attacks, Q2 2017 vs. Q2 2016

18% decrease in total DDoS attacks
17% decrease in infrastructure layer (layers 3 & 4) attacks
13% decrease in reflection-based attacks
19% increase in average number of attacks per target

DDoS attacks, Q2 2017 vs. Q1 2017

28% increase in total DDoS attacks
27% increase in infrastructure layer (layers 3 & 4) attacks
21% increase in reflection-based attacks
28% increase in average number of attacks per target

**Note: percentages are rounded to the nearest whole number*

What you need to know

- Akamai mitigated 4,051 Distributed Denial of Service (DDoS) attacks through Akamai's routed platform, a 28% increase over the previous quarter. These attacks were overwhelmingly volumetric attacks (99%).
- Egypt was the origin of the greatest number of unique IP addresses used in volumetric DDoS attacks (44,198) – 32% of the global total.
- The number of unique IP addresses used in volumetric DDoS attacks sourced from the U.S. fell 98% to 11,000 from 595,000 in the previous quarter.
- The U.S. retained the top position for both the source (112 million) and the target (218 million) of web application attacks.
- Gaming customers were targeted by 81% of all volumetric DDoS attacks. One customer suffered 558 attacks.
- Networks infected by malware using domain generation algorithms, a common command and control technique, have unique behavioral characteristics that can be used to identify them.

LETTER FROM THE EDITOR / The *Q2 2017 State of the Internet / Security Report* represents analysis and research based on data from Akamai's global infrastructure and routed DDoS solution.

The number of organizations infected and harmed by WannaCry and Petya malware gives the security community a lot to think about. We know that patching software can largely prevent damage from malware infection. And yet months after a patch became available, even after global news of WannaCry signaled a clarion call to patch, many companies still fell victim to Petya.

Patching is not a simple issue. Organizations make patching decisions based on risk and business priorities. Patching has direct costs, such as staff and testing, and indirect costs, such as downtime. Due to costs, patching is often de-prioritized as a business function. This is a legitimate decision, if it's made from a rational, risk driven viewpoint. All too often though, it's not: The conversation hasn't happened and no careful evaluation of the risks involved has been presented to business leaders.

But the risk equation is always changing. It's estimated the WannaCry malware could cost businesses \$4 billion worldwide by itself. Even the best, most rational, risk-driven decision made six months ago may no longer be appropriate today. Have any of the recent events changed the way your organization evaluates security?

This quarter's report examines trends in DDoS and web application attack traffic, along with additional research.

First, we have the DDoS Attack Spotlight, which looks at the re-emergence of PBot, decades-old PHP code that generated the largest DDoS attack of the quarter. Attackers used PBot to create a mini-DDoS botnet that launched at 75 gigabits per second (Gbps) DDoS attack.

Second, we have research showing how Akamai mined DNS-related traffic to discover anomalous behavior on networks with malware infections that use domain generation algorithms (DGAs).

Third, we have a statistical analysis of the relationship between Mirai command and control (C&C) IP addresses and their attack targets. The behavior of Mirai command and control clusters reveals that many of the individual botnets were used to attack only a few targets.

The contributors to the *State of the Internet / Security Report* include security professionals from across Akamai, including the Security Intelligence Response Team (SIRT), the Threat Research Unit, Information Security, and the Custom Analytics group.

— Martin McKeay, Senior Editor and Akamai Sr. Security Advocate

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

5	[SECTION] ¹ = EMERGING TRENDS
7	[SECTION] ² = DDoS ACTIVITY
7	2.1 / DDoS Attack Vectors
9	2.2 / DDoS Sources
9	2.3 / Industry Targets
10	2.4 / Attacks Per Target
10	2.5 / Attack Spotlight: PBOT Mini-DDoS Botnets
12	2.6 / Reflection Attacks
14	[SECTION] ³ = WEB APPLICATION ATTACK ACTIVITY
14	3.1 / Web Application Attack Vectors
15	3.2 / Top 10 Source Countries
17	3.3 / Top 10 Target Countries
18	[SECTION] ⁴ = CLOUD SECURITY RESOURCES
18	4.1 / Domain Generation Algorithm
20	4.2 / Mirai Command and Control Clusters
24	4.3 / Additional Akamai Research
25	[SECTION] ⁵ = LOOKING FORWARD



[SECTION]¹ EMERGING TRENDS

The number of IP addresses producing DDoS traffic plummeted in Q2. Of the countries that were the final hop before our network in Q1 — U.S., U.K., Germany, Canada, and Brazil — only one remained on the top five list in Q2: the U.S., where the number of IP addresses involved in volumetric DDoS attacks dropped 98% from 595,000 to 11,000. As a result, for the first time ever, Egypt topped the list of countries with the most IP addresses sourcing volumetric DDoS attacks with 44,000 source IP addresses.

While the number of attacks was up 28% after a sustained downward trend in recent quarters, the median size of attacks was reduced overall. This should not be surprising, given that our monitoring indicates that hundreds of thousands of DDoS sources were taken offline. Our research this quarter shows a botnet strain called PBot is being reused more frequently than we've seen before. This botnet has been observed with node counts in the hundreds, rather than the tens of thousands seen with Internet of Things (IoT) botnets.

Finally, for the first time in many years, Akamai observed no entries for one of the key metrics — large attacks exceeding 100 Gbps.

This quarter's Attack Spotlight on the PBot botnet reflects the trend of markedly fewer IP addresses being used in DDoS attacks. PBot node scans reveal the presence of Apache Tomcat along with the PHP interpreter. Apache Struts exploits have been observed in the wild issuing commands that attempt to download and then execute code. This is just one potential avenue that attackers are using for delivery of PBot malware. PBot botnets although limited in bot count have delivered DDoS attacks peaking up to 75 Gbps.

While we saw a precipitous drop in the number of IPs that were used in volumetric attacks this quarter, we saw a modest increase in the web application attack counts. The U.S. had the top spot as both the source and destination of the most web application attack traffic, which is a common occurrence. In fact, the attacks from most regions was relatively stable, with the exception of Asia, where attack traffic from Singapore fell by half, causing them to drop off the top 10 source country list for web application attacks.

The DDoS attacks on gaming companies certainly ramped up significantly, with one company targeted with 558 attacks over the quarter. While gaming has always been a large target for DDoS, the popularity of games relying on the millisecond timing of packets makes a tempting target, frustrating both the players and the gaming companies. This trend may culminate in significant attacks during the winter holiday season, which has often been the trend in recent years.

DDoS is a cyclic phenomenon. The chaos we've seen in the DDoS field over the past year has been monumental, and there's little reason to believe the evolution has reached a stable plateau.

[SECTION]²

DDoS ACTIVITY

After two consecutive quarters of decline in total attacks, the number of DDoS attacks increased markedly in Q2. In particular, there was a jump in attack traffic in late June. Additionally, the average number of attacks per target rose to a new high of 32.

2.1 / DDoS ATTACK VECTORS / Closer analysis of the data shows a large spike in attacks in late June. It isn't much of a stretch to surmise that this rise in attack traffic could potentially be linked to students finding ways to fill their time now that summer vacations have begun for most schools in North America, U.K. and Europe.

UDP fragment, DNS, and NTP continued as the top three DDoS attack vectors, as shown in Figure 2-1. Infrastructure-related attacks such as these accounted for 99% of DDoS traffic, as has been typical in recent quarters. Infrastructure attacks dominate because attackers find it much easier to launch a volumetric attack than an application layer DDoS attack.

Application layer DDoS attacks such as GET, PUSH, and POST floods accounted for 1% of DDoS attacks seen by Akamai. A single attack this quarter contained an SSL POST component, an attack type that is extremely rare in our experience. This was the first instance of this type of traffic in nearly two years. Most application layer attacks aren't designed for denial of service. Instead, they abuse weaknesses to breach a system.

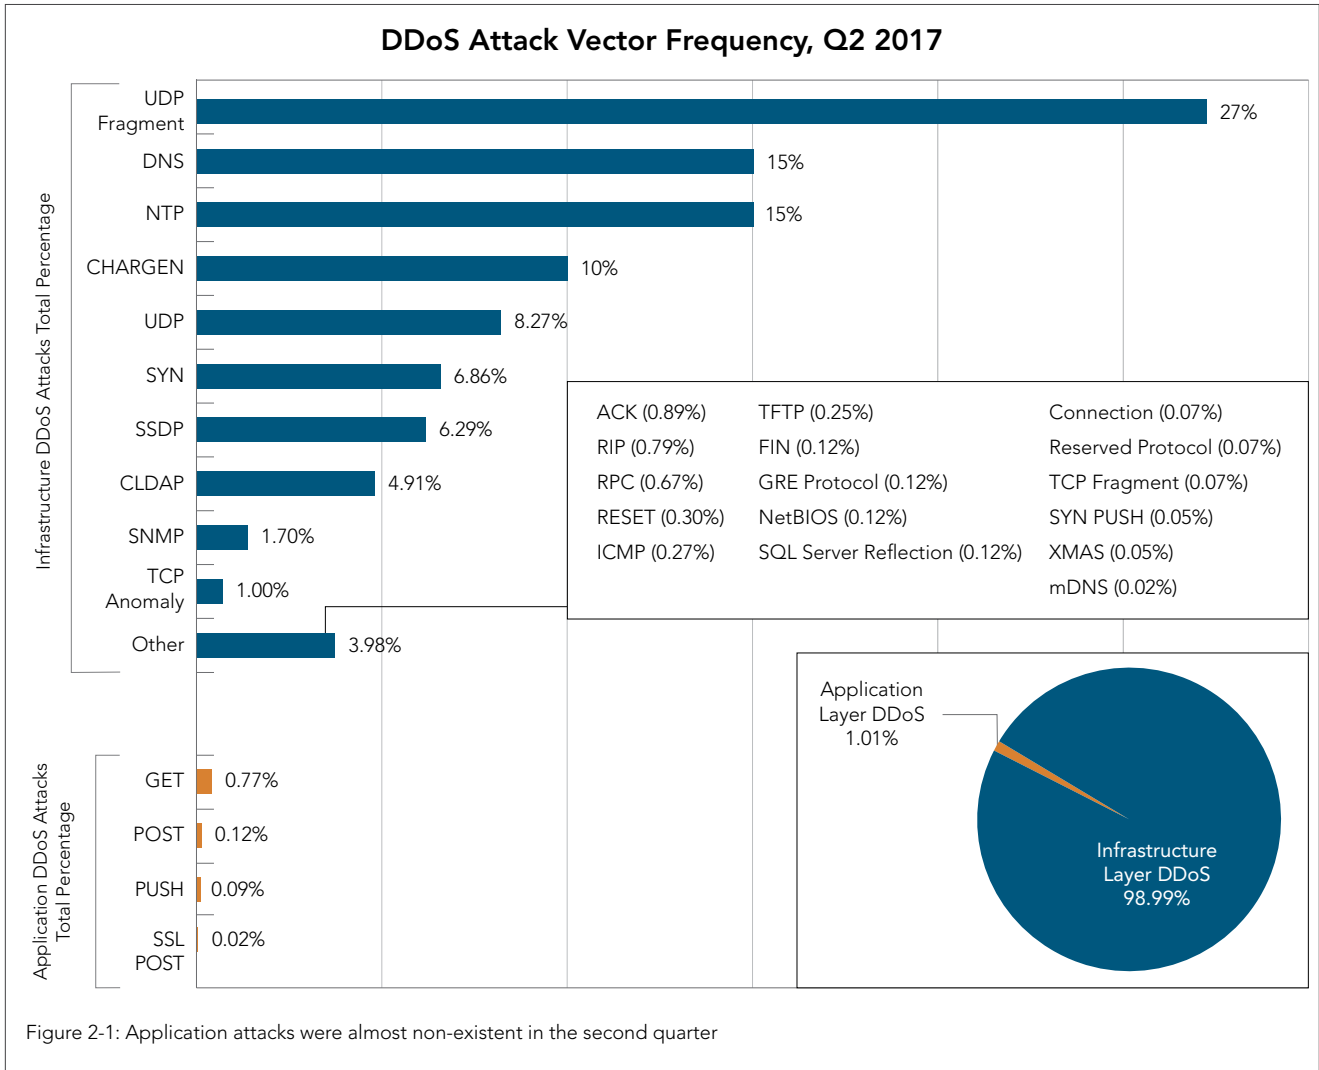


Figure 2-1: Application attacks were almost non-existent in the second quarter

As we review the 10 most frequent attack vectors per week, we see NTP, CHARGEN, and DNS continue to constitute the top three places. UDP Fragment traffic is technically in the top spot, but this is driven by the other UDP vectors and is extremely difficult to categorize. This quarter, we noticed an increase in the vector counts.

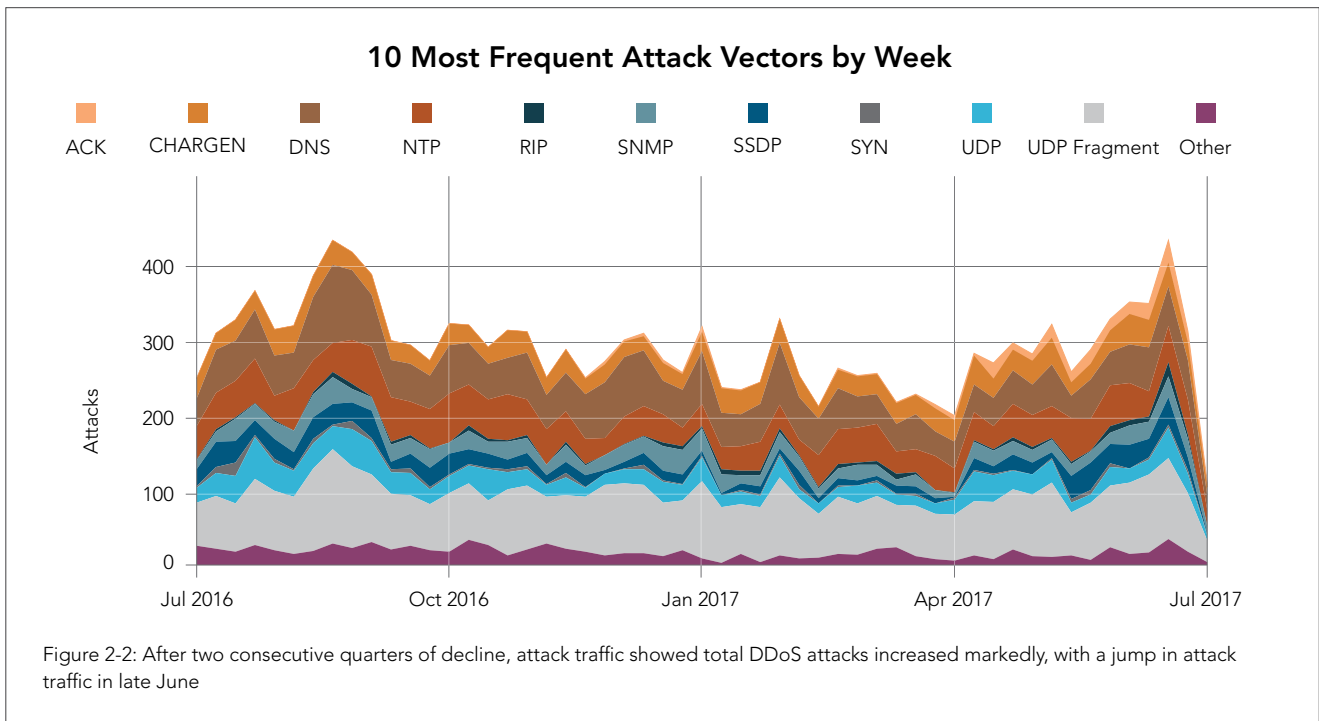


Figure 2-2: After two consecutive quarters of decline, attack traffic showed total DDoS attacks increased markedly, with a jump in attack traffic in late June

2.2 / DDoS SOURCES / Fewer source IP addresses were involved in launching volumetric attacks in Q2, as shown in Figure 2-3. The count of source IP addresses fluctuates significantly from quarter to quarter, but in Q2, it was unusually low. Egypt topped the list by producing DDoS attacks from only 44,000 IP addresses. However, in Q1, the U.S. topped the list by issuing DDoS attacks with more than 10 times that number of source IP addresses (595,000). In Q2, the U.S. was the source of DDoS attack traffic from only 11,000 IP addresses, 98% fewer than in Q1.

In contrast, during Q3 2016, China took the top spot with 81,276 IPs, while it took a count of 306,627 addresses to retain the same rank in Q2 last year.

Last quarter's unusually high number of sources may have been driven by attacks from Mirai botnets. By its nature, Mirai uses large numbers of compromised Internet of Things (IoT) devices to fuel attack traffic, each with its own IP address. Additionally, Mirai's Water Torture attack, covered in last quarter's report, creates considerable DNS traffic. In contrast to Mirai's use of small devices, PBot DDoS malware appears to have infected web servers, which can produce more DDoS traffic per device. Read more about PBot in the Attack Spotlight that follows.

Top 5 Source Countries for DDoS Attacks, Q3 2016–Q2 2017

Q2 2017		Q1 2017		Q4 2016		Q3 2016	
Country	Percentage	Country	Percentage	Country	Percentage	Country	Percentage
	Source IPs		Source IPs		Source IPs		Source IPs
Egypt	32%	U.S.	44%	U.S.	24%	China	19%
	44,198		594,986		180,652		81,276
U.S.	8%	U.K.	13%	U.K.	10%	U.S.	14%
	11,113		177,579		72,949		59,350
Turkey	5%	Germany	7%	Germany	7%	U.K.	10%
	7,049		87,780		49,408		44,460
China	4%	Canada	5%	China	6%	France	6%
	5,711		60,581		46,783		23,980
India	4%	Brazil	3%	Russia	4%	Brazil	3%
	5,224		43,863		33,211		13,502

Figure 2-3: The number of IPs involved in volumetric attacks dropped over 90% since the first quarter

2.3 / INDUSTRY TARGETS / An examination of attacks by industry vertical can provide a view into which cross section of the economy is receiving undue attention from attackers. This quarter, gaming had the lion's share of the overall attack traffic, with 81% of the DDoS attack traffic being directed at their operations. While reviewing the data, it became evident that one company in the Gaming industry was the recipient of 558 attacks during the second quarter of 2017.

DDoS Attack Frequency By Industry Q2 2017–Q1 2017

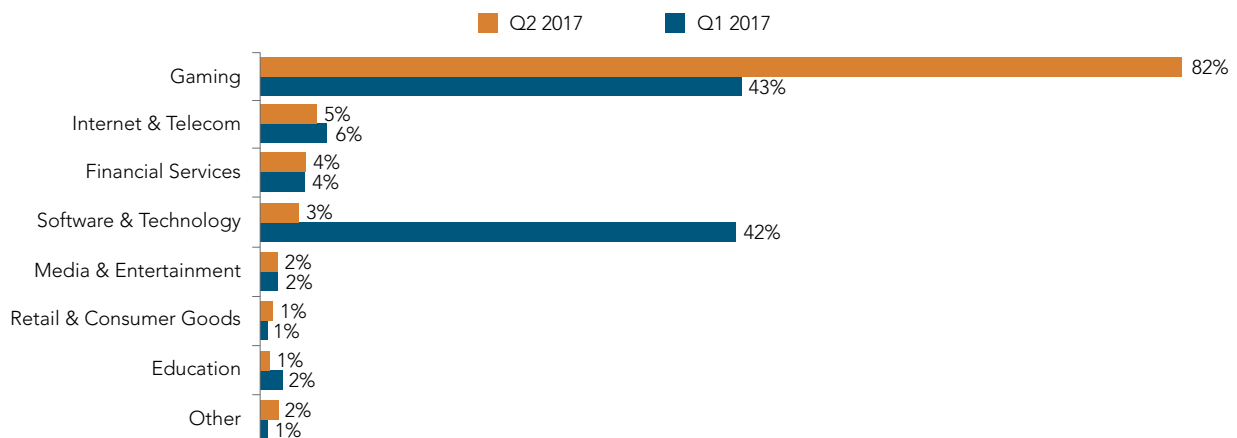


Figure 2-4: A small number of gaming sites were repeatedly hit by DDoS attacks throughout the quarter

While 81% may seem high, it is partially due to the recent reclassification of two customers from technology to gaming. As a result, some of the attack data that was previously attributed to the technology industry is now reported in the attack counts for the gaming vertical. This is a recategorization, not a fundamental change in the attack trends.

2.4 / ATTACKS PER TARGET / The overall number of attacks per target rose this quarter to a new high of 32, as shown in Figure 2-5. This average is up from 25, and beyond a more typical 30, in part due to one gaming company that was targeted 558 times - averaging six attacks per day.


2.5 / ATTACK SPOTLIGHT: PBOT MINI-DDoS BOTNETS / The current trend is smaller attacks. But attackers are also using the PBot malware to create mini-DDoS botnets capable of launching attacks with considerable power relative to the number of bots. Mirai demonstrated that weak security in IoT and other devices can be exploited to create small botnet clusters. Sources examined during a series of confirmed PBot attacks in Q2 showed potential exploitation of Apache Struts vulnerabilities. Although not as trivial to exploit as the hard-coded telnet passwords used to gain access and compromise cameras for the Mirai botnet, this vulnerability is just one example of how an attacker will leverage a weakness to gain control of a device or server.

By quickly commandeering fewer than 400 bots, an attacker can still produce enough traffic to impact a target's servers.

We know that massive DDoS attacks are possible, but could this be a new trend going forward? Have DDoS attackers taken to more subtle, targeted attacks to avoid drawing attention?

Attack Timeline and Signatures / On May 8, the first PBot attack targeted a financial customer. The attackers were most aggressive in the first few days, with the strongest attack peaking on May 9 at 75 Gbps, as shown in Figure 2-6.

Average Number of DDoS Attacks per Target, Q3 2016–Q2 2017



**TOP TARGET ORGANIZATION
DDoS ATTACK COUNT Q2 2017: 558**

Figure 2-5: Sustained attack campaigns against a few Akamai customers drove the average number of attacks per target to a new high of 32

PBot Attack Dates and Bandwidth




Figure 2-6: PBot attacks were amongst the largest DDoS attacks seen by Akamai this quarter, with the largest peaking at 75 Gbps

The attackers used PBot's UDP flood attack command to produce these attacks.

```
PBot UDP flood - Port 80
18:17:42.952535 IP X.X.X.X.219.43424 > X.X.X.X.80: UDP, length 1400
18:17:42.952537 IP X.X.X.X.41065 > X.X.X.X.80: UDP, length 1400
18:17:42.952539 IP X.X.X.X.219.43424 > X.X.X.X.80: UDP, length 1400
E...W0@.7...s.I..&s
...P..|.G.(|...wV.....z?.3./|.6....A.Q}.....L....dM/.    %3D..$....h.<E....}.M.
.tA.|8...;.@l.G).^...~pzE.?O..Dv.}.B.\...e.Qh:.....EF^l..bqx[U...g\Y.T....3Q...6.&3...g/
{^dI.....5ls...m..["...R.~|v..#]..&.../.R..D.dg9.....~$......
.4..H...G.g...y.t...n.n./9...w.P.$(F>.-.Kt.!.9J.'...~.....Y.....
<snip>

PBot UDP flood - Port 53(DNS)
09:53:19.272125 IP (tos 0x8, ttl 116, id 11246, offset 0, flags [none], proto UDP (17), length
1428)
  X.X.X.X.51818 > X.X.X.X.53: 31094 op6% [b2&3=0x3090] [36437a] [48580q] [11198n] [62065au]
Type13021 (Class 25108)? [|domain]
09:53:19.272129 IP (tos 0x0, ttl 48, id 7724, offset 0, flags [DF], proto UDP (17), length 1428)
  X.X.X.X.56792 > X.X.X.X.53: 60713% [b2&3=0x67d] [36280a] [56979q] [48520n] [23417au]
Type13990 (Class 57257)? [|domain]
09:53:19.272131 IP (tos 0x8, ttl 116, id 11247, offset 0, flags [none], proto UDP (17), length
1428)
  X.X.X.X.51818 > X.X.X.X.53: 31094 op6% [b2&3=0x3090] [36437a] [48580q] [11198n] [62065au]
Type13021 (Class 25108)? [|domain]
E...+...t.C.....&r..j.5..{uyv0...U+..q..2.b....`...8....=.d.....#+.....}.!...q.....
kZF="c.....5..lY...J..o.O..#...Yq..a3lv..tg.....3n.....:ad...X!.v..x.z%T....7..
[...@r...x.1.36BxXUjD"...F.b.D.....L.f..w..V<.....i.....|.....P..A.....
.^.....<Y.O.....x..Y.....IXC.....g..>.x.l.....[.....>.&..C..
RGNY@..}.....`. <R/.qk.l..i..u.p..).y<..... .wx...Y...G...;..8.Bc.....& .E....V-.....l\p.
<snip>
```

Figure 2-7: Signatures from two PBot UDP floods, one targeting port 80 and the other targeting port 53

The signature from the first attack targeted port 80 over UDP. Most likely this attack was an attempt to consume bandwidth or disrupt users of the site/service. The attacks on May 9 were set to target the DNS service on UDP port 53. The attacks were all UDP floods however, some packet analysis tools, such as tcpdump, will attempt to decode packets based on target port. As a result, the random UDP payload sent to port 53 is displayed as DNS even though it is not. The attackers may have conducted reconnaissance of the target network or had knowledge of available services.

PBot IRC Attack Tool / Attackers sometimes recycle old attack tools and scripts capable of producing DDoS attacks, which is the case in this series of attacks. Scans of infected bots revealed that they could execute PHP code. Scans of infected bots also revealed that besides running PHP interpreters, they were also running Apache Tomcat. Based on recent Apache Struts vulnerabilities, this provides a potential, yet unconfirmed, means to deliver and execute the PBot code. Akamai SIRT obtained a sample of the PBot attack script for dissection and testing in a lab environment. The signatures of the UDP flood matched exactly with attacks observed in the wild.

```
15:46:36.167200 IP X.X.X.X.37777 > X.X.X.X.55: UDP, length 1400
15:46:36.167205 IP X.X.X.X.37777 > X.X.X.X.55: UDP, length 1400
15:46:36.167209 IP X.X.X.X.37777 > X.X.X.X.55: UDP, length 1400
E...S@.~@.....7...G...T.@W...bPn.....%}V..H.!p.%If.]v.#.\.../...+.{.t..0...k
.2.....n.l.n..4..j..._2..%l...k.|...UJ.....6T...>.p.7..C.....0...|.K...k... ..D.G.B..
Bm.M..$......Zv..a..b...N.....w..#.....Q.....s..UV.{6.....`.r.W..OR.~.z.M...c.[1
VlyM^(..Q..`A)?.?...Ll...5...}2..v.+..C].s.k.....-A.$f....?..?..>..+6.?..!..M.X.;fl.....S.
.F..XN..E.p...*.M..4y.|R.Ek.....G.....'3.bF..6.K.f.rN.....X.A.z..J,\...#SY....T..".I.+.....)
x0.{}.Q.5..Q&...;.p.z...W..q...{.q.H@M`u/_...>.D\...w?~2\b.._..m.Ba.3Iqcbm
```

Figure 2-8: This attack signature indicates it's often easier to recycle code than to create new tools for each botnet

PBot controls its bots through Internet Relay Chat (IRC). This is a classic command-and-control architecture. The attacker sets up an IRC server for the bots to connect to. Once a potential bot runs the malicious PBot PHP code, it will connect into the specified IRC channel. The settings for this communication are included within the source code.

It's trivial for malicious actors to modify the Pbot code for their own use. As of this writing, Akamai SIRT has observed three versions of the PBot code used in DDoS attacks. Once a bot connects, it's ready for commands that include carrying out attacks. An example of the -UDP command prompt and response is shown in Figure 2-9.

```
16:02 -!- Irssi: #<channelname>: Total of 2 nicks [1 ops, 0 halfops, 0 voices, 1 normal]
16:02 -!- Irssi: Join to #<channelname> was synced in 1 secs
10:59 < blackbox> -udp X.X.X.X 55 30 1400
10:59 <[<>]BotID> [Attack Sent! X.X.X.X / 30]
```

Figure 2-9: PBot attack command parameters -UDP IP port (55) time (30) packetsize (1400)

Conclusion / PBot is an example of smaller DDoS attacks this quarter. Attacks have peaked at less than 100 Gbps, but can still have an impact. It's important to be aware of services available within a network and potential avenues of attack, particularly which ports are open and accessible to the Internet. With large DDoS attacks, it's possible for an attacker to completely consume all available bandwidth at the target site. For these smaller DDoS attacks observed from PBot and other bots, it has been common for specific ports to be targeted. The goal of the attacker is to have as much impact as possible while working with limited resources.

The PBot code is PHP-based, which also allows for multi-platform operation. The same services that can be targeted for attack can be potential avenues for exploitation. Take steps to avoid becoming part of the PBot problem. This could include making sure patches are applied in a timely manner or unnecessary services are firewalled to avoid access from external sources.

2.6 / REFLECTION ATTACKS / Reflection attacks continued to dominate DDoS activity in Q2. As noted in prior quarters, we see that DNS, NTP, and CHARGEN maintained their spot as the top three reflection attack vectors. Their ongoing use by attackers demonstrates a collective need to address this problem. This can be best addressed by system administrators applying security patches and ensuring that systems are configured with security in mind. DNS configurations should be reviewed on a regular basis or be assigned to a vendor partner.

Reflection-Based DDoS Attacks, Q2 2016–Q2 2017


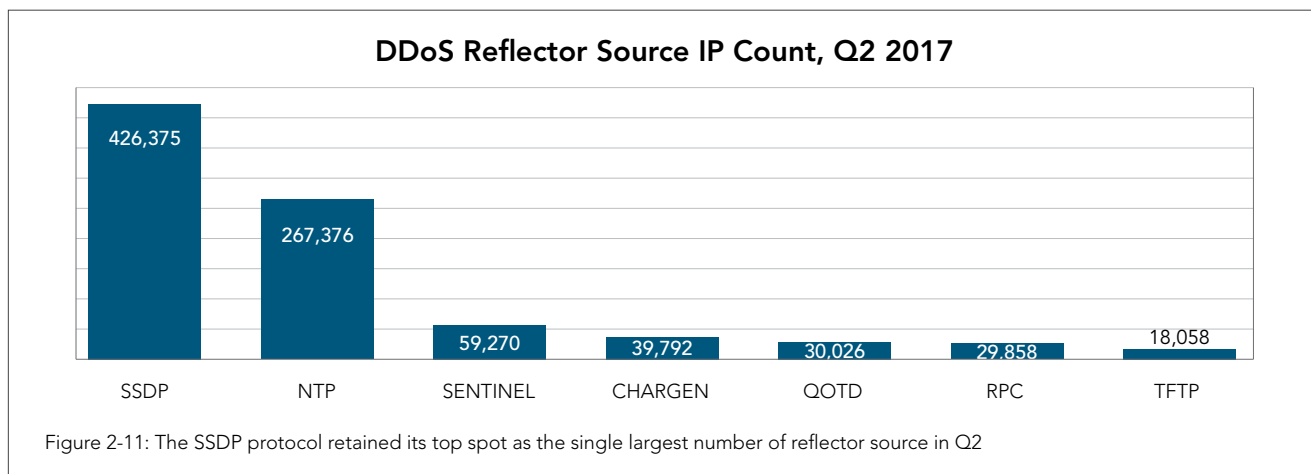
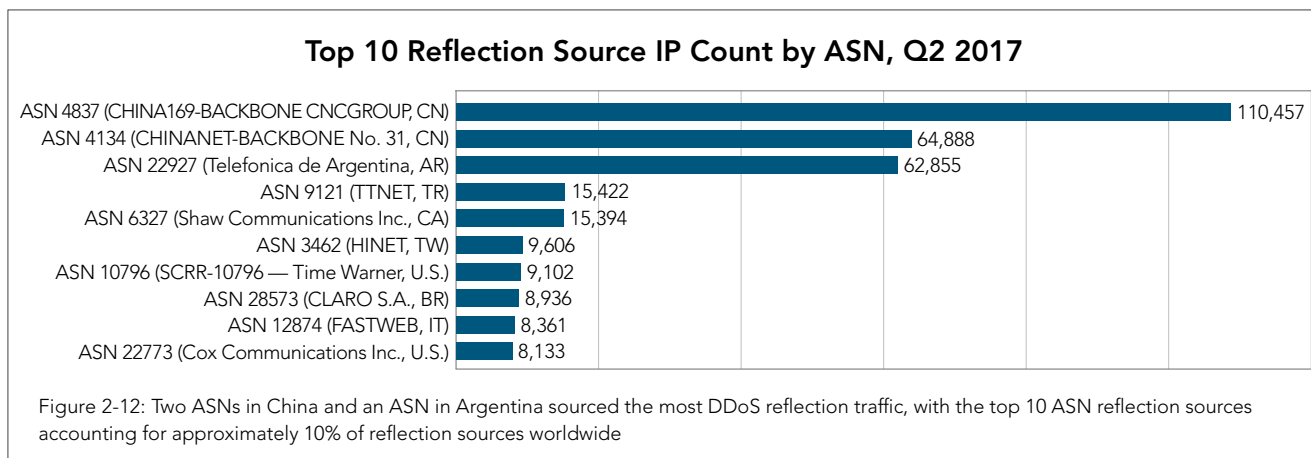


Figure 2-10: Total reflector count grew this quarter, but are still lower than at the same time last year

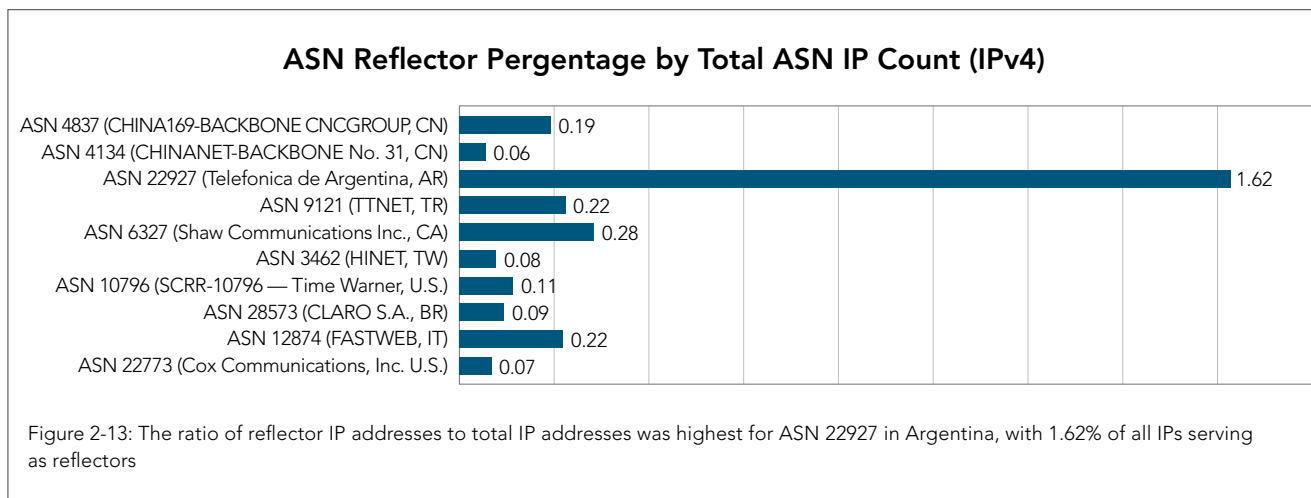
Of all DDoS reflection attacks in Q2, 33% used DNS reflectors attacks, 28% used NTP reflectors, 17% used CHARGEN reflectors, and 12% used SSDP reflectors. Overall reflector count across all vectors is lower than at the same time last year.



Two ASNs in China and an ASN in Argentina sourced the most DDoS reflection traffic, as shown in Figure 2-12. The makeup of this list has changed very little in recent quarters, though the number of reflectors in each ISP has been growing over time.



The ratio of reflector IP addresses to total IP addresses was highest for ASN 22927 in Argentina, with 1.62% of all IP addresses serving as DDoS reflectors. It may not seem like a high proportion of the IP addresses owned by Telefonica de Argentina are reflectors. But when this ASN is compared to any other member of the list, it becomes apparent that it is much easier to find a usable reflector in this network than any other we're tracking.



[SECTION]³

WEB APPLICATION ATTACK ACTIVITY

Application layer attacks continued to slowly grow with a 5% increase quarter-over-quarter and a 28% increase year-over-year. Unlike DDoS attacks, web application attacks involve relatively little traffic and can be hard to detect, which in many ways makes them more dangerous.

3.1 / WEB APPLICATION ATTACK VECTORS / SQL injection (SQLi) attacks were used in more than half (51%) of the attacks, nearly 185 million alerts in the second quarter alone, as seen in Figure 3-1. This is up from 44% of all attacks in the first quarter. Attackers know these vulnerabilities exist in many sites and put increasing resources into finding ways to compromise them. These attacks are automated and look for any vulnerable system, rather than target specific organizations.

Local File Inclusion (LFI) was the second most used attack vector of the quarter (33%/121 million alerts), followed by XSS (9%/33 million), RFI (2%/8.6 million), and PHPi (2%/6.1 million). While dwarfed by other attack types, Java injection attacks have grown by 800% since Q2 2016.

Web Application Attack Frequency, Q2 2017

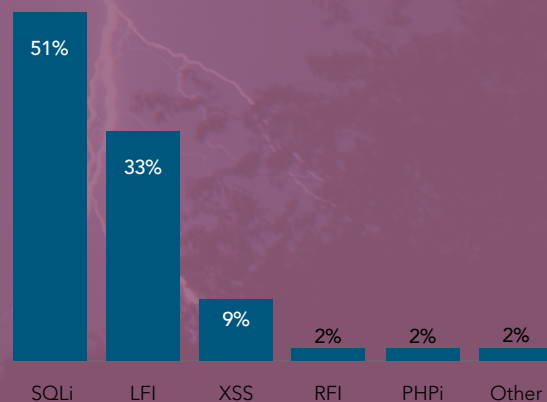
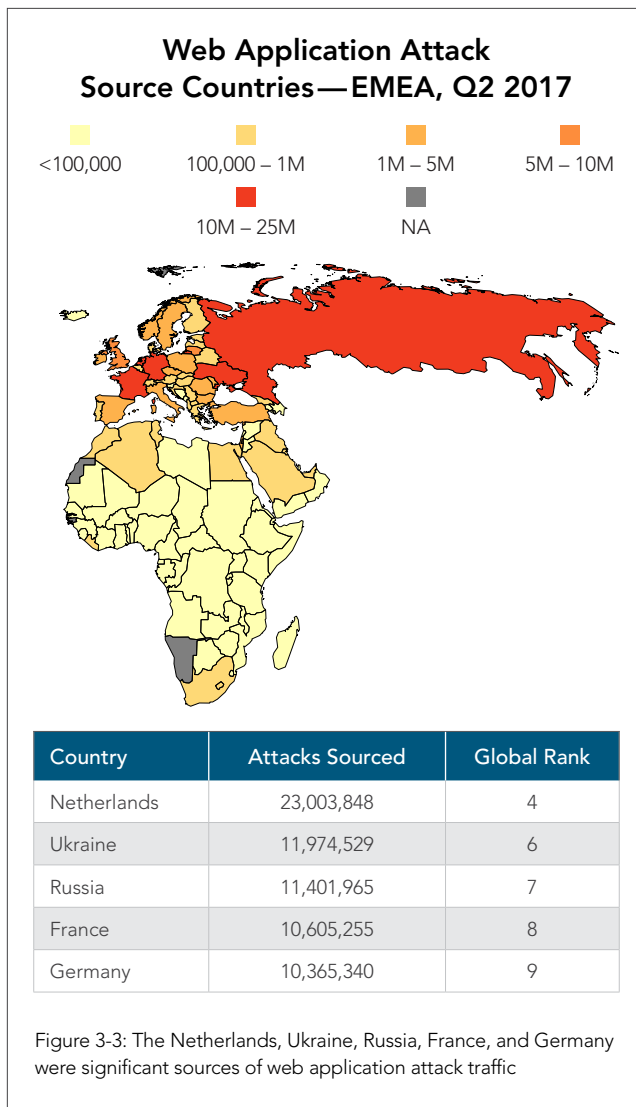
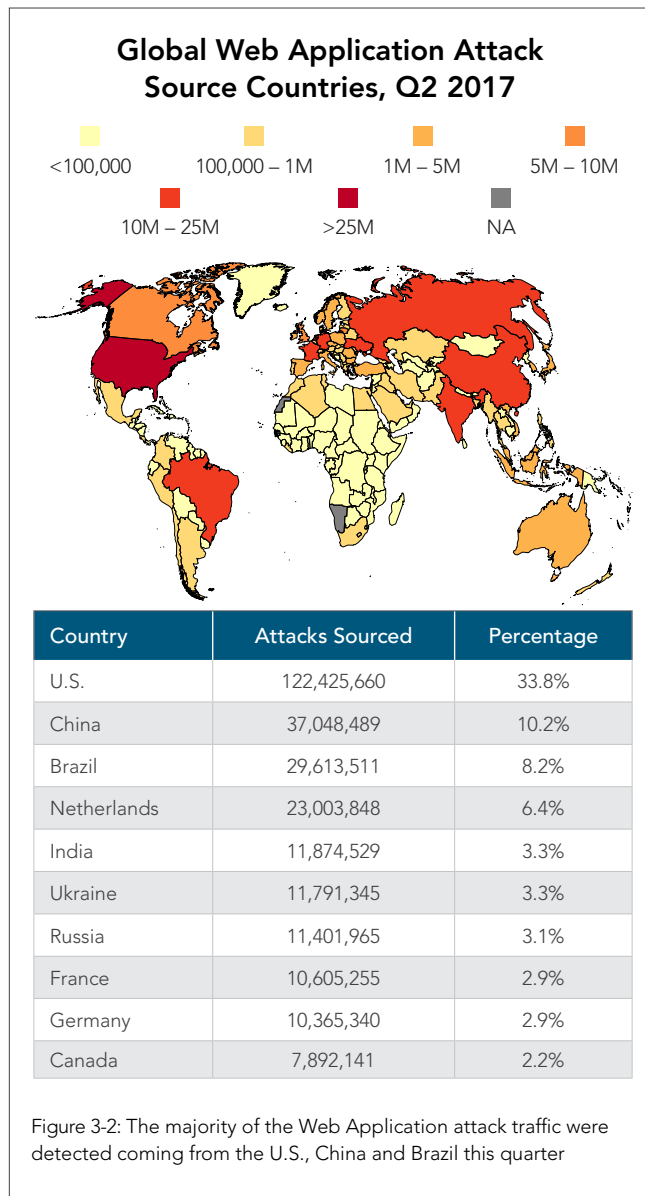


Figure 3-1: SQLi, LFI and XSS attacks accounted for 93% of web application attacks in Q2

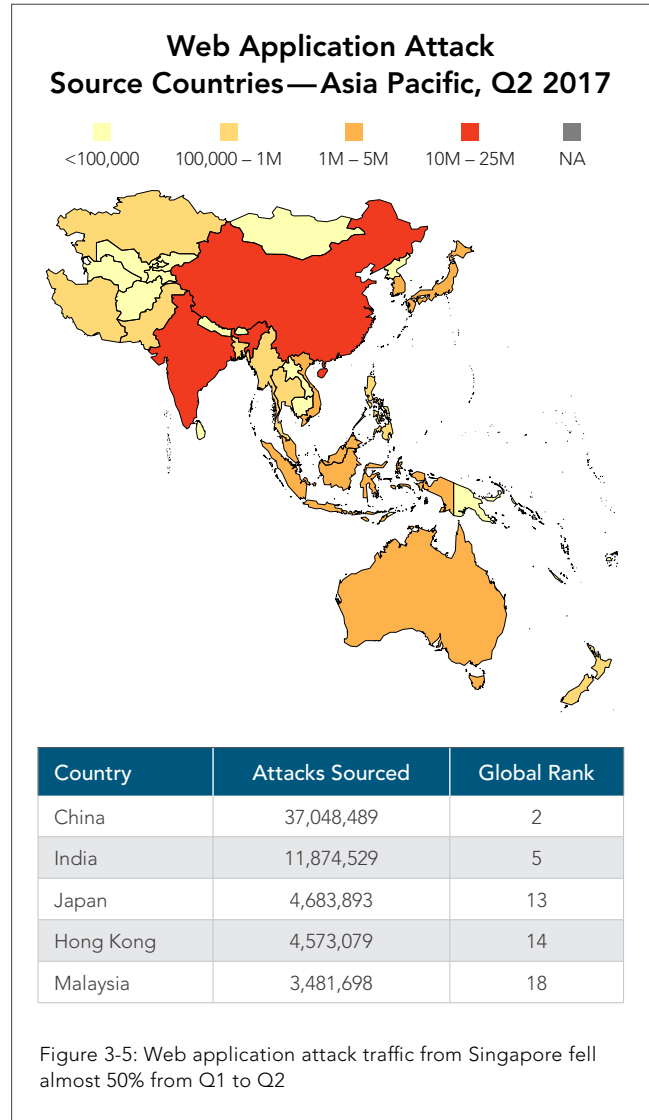

3.2 / TOP 10 SOURCE COUNTRIES / In Q2, the top source countries for web application attacks remained relatively unchanged with one key differentiator: Canada was one of the top ten source countries. Canada had been in 11th place for the last two quarters. The top three origins of attack traffic were the U.S. (33%), China (10%), and Brazil (8%).

In Europe, the Middle East and Africa, detected attacks fell drastically in the second quarter. The Netherlands is still responsible for the most significant portion of alerts, despite falling from nearly 44 million attacks to slightly over 23 million. One issue that arose during this quarter is spoofing of traffic using X-Forwarding-For headers to make the traffic appear as if it was from another country. This is a type of spoofing that is unusual for us to see and merits further research.




The attack landscape in the western hemisphere grew slightly between the beginning of April and the end of June. This quarter, Canada moved into the third position overall in the Americas. There's a noticeable gap between the U.S., Brazil and Canada, with millions of detected attacks per country and Mexico and Peru, where our measurements were in the hundreds thousands instead.

As illustrated by Figure 3-5, China remained the overall top source country for web application attacks in the Asia-Pacific region, followed by India and Japan. Noticeably absent from the top five was Singapore, which was fourth place overall in Q1.



3.3 / TOP 10 TARGET COUNTRIES / The U.S. maintained its position as the largest target of attack traffic, with more than 218 million attack triggers. With a 130% increase in attacks since Q1, the U.K. was the second most targeted country in Q2. Brazil saw a 15% increase in attacks since last quarter, placing it third. China returned to the list in ninth position, while Singapore moved to fifth place, after being the tenth most frequent target in Q1.





[SECTION]⁴

CLOUD SECURITY RESOURCES

4.1 / DOMAIN GENERATION ALGORITHM / DNS plays a significant role in the interconnected world of malware-infected devices communicating with their command and control (C&C) servers. DNS activity on infected networks is different from DNS activity on clean networks. By identifying the relevant characteristics and practicing machine learning algorithms, we can identify abnormal behaviors that lead us toward detecting malware activity.

The main purpose of the DNS protocol is to act as a translation layer between server/application name to the IP address. It creates a many-to-many relationship between computing resources and application names. DNS was intentionally designed as a protocol to allow for easy and frequent re-assignment of domain names to different computers, services and devices, regardless of hosting platform or country. These same capabilities, when exploited by malicious C&C servers, can help them remain almost invisible while defenders try to track them and take them down.

A more sophisticated technique is being used by malware developers to build scalable malware botnets that persist longer. This technique is being used by the majority of modern malware and is known as *Domain Generation Algorithms (DGA)*.

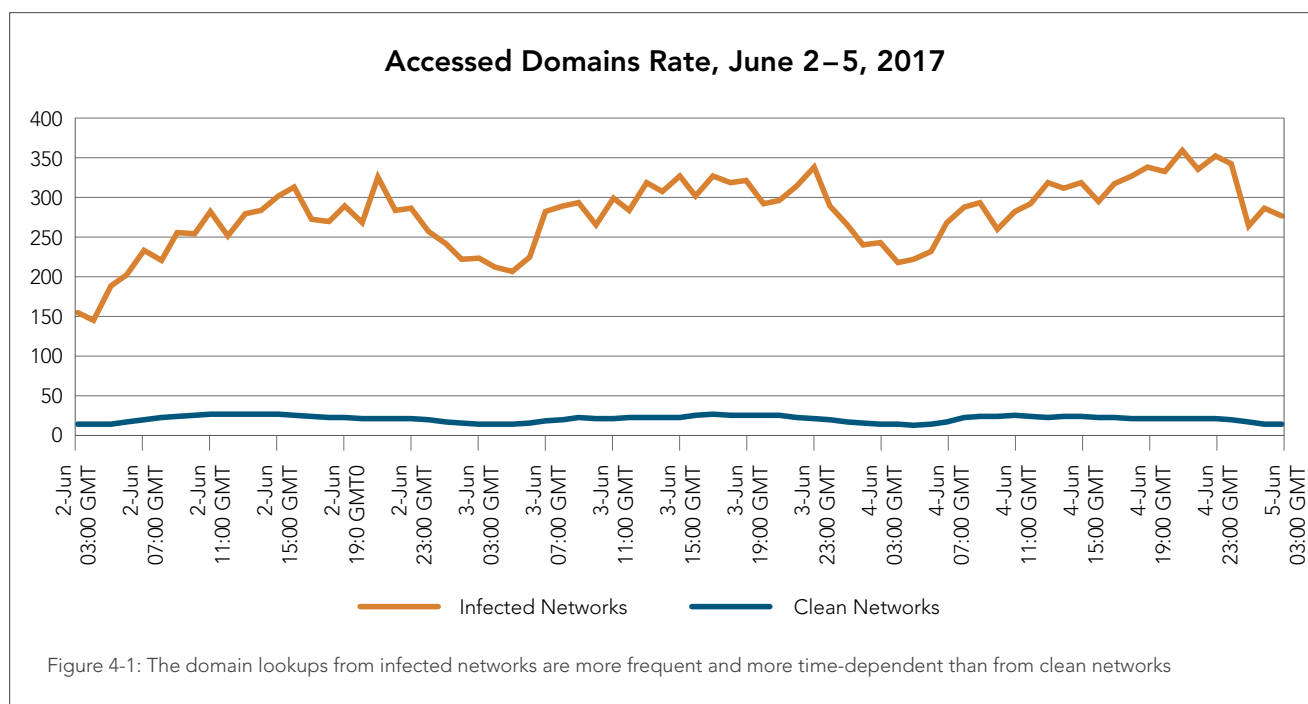
The DGA communication technique was first introduced with the emergence of the notorious Conficker worm in 2008. The first variant of Conficker (variant A) generated 250 different domain names each day, using the date as a seed, allowing generation of the same random domain names across all malware instances every day. The FBI responded by using a reverse-engineered piece of Conficker to register all(!) the domains the malware was going to use before the malware operators had the chance to do so themselves. As the Conficker malware evolved, introducing another 4 malware variants (B to E), the usage of randomization techniques also evolved. Variant C generated more than 50,000 domains per day out of which only 500 would be randomly chosen to try to communicate with the C&C server. This made the prevention step taken by the FBI far more expensive.

The ability to generate an endless amount of random domain names daily makes the work of taking over malicious domains nearly impossible for defenders. Once the bad actors that developed the malware want to control it, they only need to register one of the domain names expected to be generated. At which point, botnet members will start communication and authentication with that newly activated C&C server.

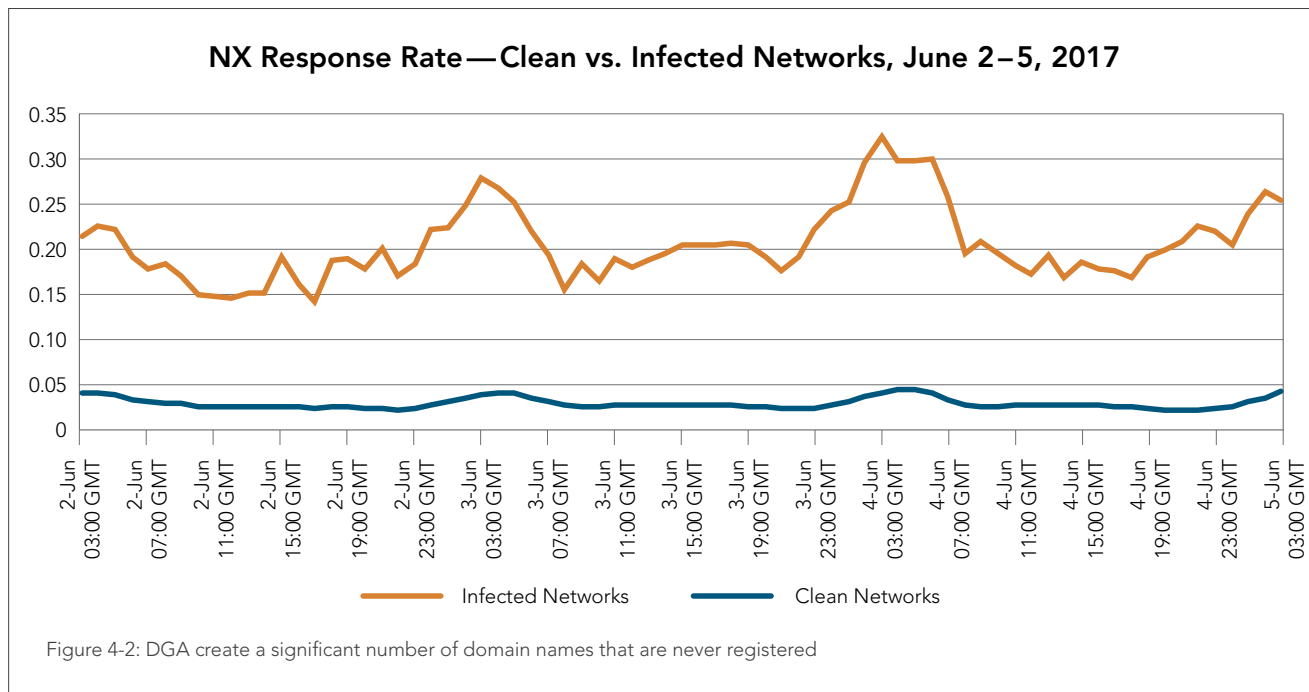
Akamai's Enterprise Security Research Team performed an in-depth analysis of the behavior of network traffic with and without the presence of DGA-based malware.

As part of the research, a sample of U.S.-based service provider DNS traffic was inspected for 48 hours. Here are some statistics and graphs that show traffic characteristics over time:

- The majority of the networks were small to medium networks representing, in many cases, home routers.
- The sampled traffic contained between 2.5 to 2.75 million connected networks at any given moment.
- There were 140 networks infected with malware in this analysis.



When looking at the average number of unique domains accessed per hour, we saw that infected networks had approximately 15 times the lookup rate of a clean network. This can be explained as the outcome of access to randomly generated domains by the malware on the infected networks. Since most of the generated domains were not registered, trying to access all of them created a lot of noise.



When looking at the percentage of DNS NXDomain response codes (this response code implies a domain does not exist), we can see that clean networks have 1%–5% NXDomain responses, while infected networks have 15%–33% NXDomain responses. These numbers can be explained as the outcome of malware looking up random domains that don't exist, causing DNS servers to respond with NXDomain.

We can see clearly that infected networks have a different set of behavioral characteristics. Once we start looking at these relevant characteristics and practicing machine learning algorithms, we can identify abnormal behaviors that lead us toward detecting malware activity.

Summary / While it has been a nearly a decade since the first time DGA was used by a malware, it remains a frequently used communication technique for today's malware. Over the years, DGA techniques evolved and improved to empower malware to become robust and resilient, making it more resistant to takedowns and other defender actions.

Tracking networks' DNS traffic and determining whether communication is being used for malicious purposes is a challenging task. It requires observation of a landscape of normal network activity, which can lead to the understanding of the abnormal activity we are all looking for. When it comes to Internet connectivity, DNS traffic plays a significant role in security monitoring. It requires defenders to have visibility into current blind spots that may be overlooked, such as roaming users and IoT devices.

DNS is frequently used as a component of the communication between malware and C&C server and is deserving of further investigation in most organizations. DNS traffic can be monitored to discover infections that might otherwise evade detection.

Security defenders are advised to make certain they are using a combination of security monitoring products that include DNS monitoring. Having visibility into different areas of the enterprise network will increase detection and reduce risks. For the best defense, security controls should also be in place on endpoint devices and the inner network, not just Internet connectivity.

4.2 / MIRAI COMMAND AND CONTROL CLUSTERS / When the Mirai botnet was discovered last September, Akamai was one of its first targets and we continue to receive attacks from the botnet to this day. So it shouldn't be surprising that our researchers have continued studying different aspects of the botnet.

The command and control (C&C) structure of Mirai is of particular interest to Akamai. Our initial research examined more than nine months (288 days) of the collected data in order to better understand the nature of the shifting structure that issues commands to the bots that actually carry out the attacks. In other words, this is preliminary research on the C&C's, not the end nodes, or bots, that send traffic at the target.

We often talk about Mirai as if it were one large network of bots, but in reality, it is more akin to smaller hives of related bots and C&Cs. We define each individual compromised system that sends attack traffic as a bot or end node. The term C&C primarily refers to a single system with a single IP address, though it can also be used as a general term for all such systems. A cluster refers to a group of related C&Cs that are connected by DNS information over time and are assumed to be controlling the same set of bots. We call a group of bots that respond to a C&C cluster a botnet.

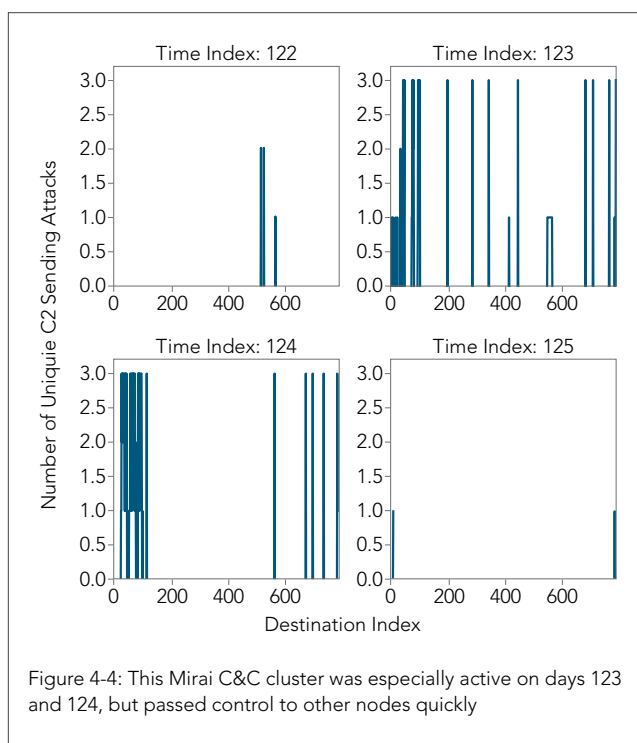
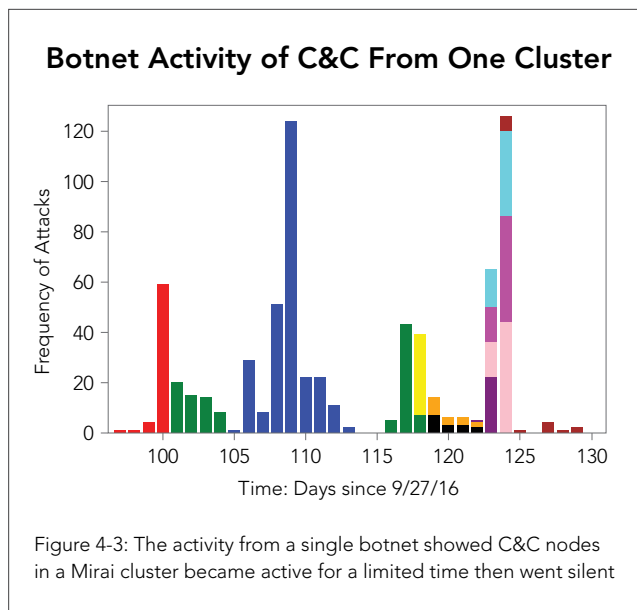
Understanding what makes a cluster of C&Cs requires significant effort and historical data. Starting with an IP of a known C&C server, we made a connection to the DNS history associated with that IP. Once the DNS relationship was established, we looked at the other IP addresses that had been associated with the domain and added those to the cluster. The second group of IP addresses were then reviewed to see what additional DNS records were associated. This process was repeated until all linkages were followed. The resultant set of IP addresses was designated as a cluster. This was further validated by reviewing the targets the botnet controlled by the cluster attacked, and we found very strong correlation between these IPs.

In Figure 4-3, we've chosen a single cluster, with 32 days of activity, as a representative example of the botnet's C&C structure. This plot shows the activity for 12 of 24 C&C nodes, limiting the data to the first nodes that were seen. Each color represents a unique IP address that was used during the time period. We chose to limit the number of IPs in the plot in order to clarify the example. The height of each line shows how many attack commands the C&C cluster sent on a particular day.

We saw a common thread emerging at the start of the time series: A single C&C node was active for a limited number of days, then died off, not to be seen again. The green node was slightly different in that it was inactive for more than 10 days before becoming active again. This specific cluster of C&C nodes was especially active on days 109 and 124. On day 124, the cluster used four nodes to issue commands, compared to the one node used on day 109, during the first attacks.

When we look at the same node from the target viewpoint, we see the attack commands as a grouping of destinations. Each line in Figure 4-4 represents a separate target IP as a destination index. In other words, the destination is assigned a number that has no other context, and proximity does not show any relationship amongst the target IPs. The height of the line indicates the number of C&C nodes that issued attack commands.

Comparing the two figures, we can see that the cluster of C&C nodes was especially active on days 123 and 124. While the plot shows us that there are a number of targets that were attacked almost continuously by the botnet, there was also a series of transitory targets that only showed up for a limited time frame with any C&C cluster. In the future, we plan to investigate similar attack patterns on the targets using shorter time scales — hours or minutes instead of days — to more accurately understand the cluster structure of the botnets.



The cloud of dots shown in Figure 4-5 shows all networks targeted by Mirai nodes during our monitoring. The size of the dot is indicative of the number of attack commands issued against that target. Akamai is shown in orange, outlined in black. This plot shows that there was a moderate number of targets hit almost continuously, and a cloud of targets that were only hit a few times during our data collection. For comparison, the largest dot represents more than 10,500 attack commands during our monitoring, while Akamai was targeted 1246 times. We apologize for creating a plot that's difficult to read if you are colorblind. It might be difficult to read even if you're not.




Figure 4-5: Akamai (Orange with a black outline) was a major target, but some had it worse

A few interesting points emerge when we look at the targets for each C&C cluster, shown in Figure 4-6. The top left facet shows the same aggregate data as Figure 4-5, while all other facets show the commands of a single C&C cluster. The dots indicate the number targets the cluster received commands to attack. If there is only one dot, there was only one target. If there are multiple dots in a facet, there were multiple targets and the size of the dots show the number of attacks a target received as a proportion of the whole for that cluster.

More than a few clusters concentrated on a single target for their entire lifespan, while the majority of the clusters were used to target a diverse range of endpoints. The clusters that only attacked a small number of targets will make for interesting analysis to see what sort of relationships emerge. At least one botnet operator was offering access to the systems under its control for rent, which may explain why some botnets attacked such a large number of IP addresses.

This research is based solely on the data that was available to our researchers and, if anything, underrepresents the entirety of the Mirai botnet. We are certain that there are attack commands that were not seen by Akamai's research team, but we also have very high confidence in the data we've represented by this research.




Figure 4-6: Akamai (in orange) was targeted by 42% of the Mirai C&C clusters

4.3 / ADDITIONAL AKAMAI RESEARCH /

Passive HTTP2 Client Fingerprinting— White Paper— Akamai’s Threat Research team recently conducted research on the possibility of passively fingerprinting HTTP2 clients based on unique implementation features. The paper also proposes a format for passive HTTP2 fingerprints, as well as a few examples of unique fingerprints belonging to common clients and implementations.

DDoS Attacks against DNS Infrastructure in the News— DNS infrastructure is a ripe target for malicious actors hoping to disrupt a digital property’s availability because it provides the initial resolution for an end user’s browser client from hostname to IP address. At best, an attack against your DNS records can significantly delay an end user’s connection. At worst, it can render your application inaccessible to the end user, either through a denial of service or through a DNS record hijack or forgery.

Low Risk Threat: DDoS Extortion Letters— Adversaries calling themselves the Lizard Squad have been sending businesses extortion letters, demanding payment in bitcoin to prevent a Distributed Denial of Service (DDoS) or other attack against their applications. These letters have been sent to businesses across the globe and across industries for several years, with little follow-through.

Spotlight on Malware DGA Communication Techniques (See also [4.1: Domain Generation Algorithm](#))— This well-known technique empowers malware developers to build scalable malware botnets that live longer. In fact, it’s being used by the majority of modern malware and is known as Domain Generation Algorithms (DGA). In this article, we will tell the story of DGA. We explain when it was first introduced. We also discuss how it’s being used in the wild and what challenges defenders face, and finally we address how we can fight back using machine learning and behavioral algorithms.

WannaCry: What We Know— On Friday, May 12, news agencies around the world reported that a new ransomware threat was spreading rapidly. Akamai’s incident response teams and researchers worked quickly to understand this new threat and how to mitigate it.

Dealing with Petya— Akamai is aware of and is tracking the malware threat known as “Petya.” Petya is ransomware spread using several methods, including PSexec, Windows Management Instrumentation Command-line (WMIC), and the EternalBlue exploit used by the WannaCry family of ransomware.



[SECTION]⁵ LOOKING FORWARD

Have you started your Christmas shopping yet? Neither have we. But we have started to gear up for the holiday season. If you've ever worked at a merchant with an online presence, you are probably aware that there comes a day in the late third quarter or early fourth where all changes, all patches, and all new projects stop in preparation for the holiday shopping season. Then, some time after the beginning of the new year, all the pent-up updates and changes will be let loose in one large flood. This is sometimes known as the Holiday Freeze.

It may seem a little odd to be asking about Christmas when the year is little more than half over. Planning for future threats and risks is as important to security organizations as planning for the holiday shopping season is to our counterparts in marketing and sales departments of retail organizations. For most organizations, security events aren't seasonal, they happen year-round, without the ability to anticipate attacks. Unless you're the security team for a merchant,

in which case you need to plan for Black Friday and Cyber Monday, since they are likely to be the high water marks for attack traffic for the year.

This quarter saw a rise in the number of both DDoS and web application attacks targeting organizations that Akamai protects. The numbers show a rebound in DDoS attacks compared to the last several quarters, when their frequency was slowly diminishing. Does this indicate that we're going to see future increases in the number of attacks? There's no way to be certain, but we do know that both web application and DDoS attacks are cyclical, and that they often return, more powerful than ever. It would require a tectonic shift to the nature of the Internet to change this truism. So, like planning for the holiday season, we have to plan for the next high tide of attack traffic.

Speaking of planning, repeat readers of the *State of the Internet / Security Report* may notice we're making a change to the type of content we report on. This report has historically concentrated on DDoS and Web Application attacks, which we will continue to examine in depth each quarter. But as a company, Akamai has multiple projects and datasets related to security, which have not been explored in full. In the Domain Generation Algorithm story and the Mirai Command and Control Clusters research, we wanted to give you a taste of the type of research you'll be seeing in the future.

Two of the stories we have planned for the coming quarter are a blog post on Fast Flux command and control structures for botnets and a story on DNS exfiltration. Fast Flux is a method used by botnet handlers to hide as much of their traffic and command infrastructure from defenders as possible. In contrast, DNS exfiltration is used to hide data being pulled from your network using DNS requests that are often ignored by many organizations. This quarter's research into the Mirai is a precursor to further intelligence.

Whether it's Hanukkah, Christmas or simply time off at the end of the year, we need to be planning for the future rather than reacting. The majority of security risks are understood, even though constant variations and fluctuations force re-evaluation of what the top risk is at any given point. The more we can think about the future, about how those trends are going to manifest over the long term rather than reacting to individual incidents, the better we are able to put in controls that are going to protect us today, tomorrow and into the future.

STATE OF THE INTERNET / SECURITY TEAM

Jose Arteaga, Akamai SIRT Lead, Data Wrangler — Attack Spotlight
Dave Lewis, Global Security Advocate — DDoS Activity, Web Application Attack Activity
Chad Seaman, Akamai SIRT — Attack Spotlight, Mirai Command and Control Clusters
Wilber Mejia, Akamai SIRT — Attack Spotlight
Alexandre Laplume, Akamai SIRT — Attack Spotlight
Elad Shuster, Security Data Analyst, Threat Research Unit
Or Katz, Principal Lead & Security Researcher — Domain Generation Algorithm
Jon Thompson, Custom Analytics
Shrijita Bhattacharya, Intern — Mirai Command and Control Clusters

EDITORIAL STAFF

Martin McKeay, Sr. Security Advocate, Senior Editor
Amanda Fakhreddine, Sr. Technical Writer, Editor

EXECUTIVE ADVISORY GROUP

Alex Caro, Cliff Crocker, David Duff, Charlie Gero, James Kretchmar, Mick Scully, Josh Shaul, John Summers, Dan Walter, Kate Zinn

INTERN SPOTLIGHT: SHRIJITA BHATTACHARYA

Shrijita Bhattacharya has been the SIRT Summer Intern for summer 2017. She is currently a fourth year PhD student in the statistics department at University of Michigan, Ann Arbor. She worked alongside Chad Seaman for this report, looking closely at Mirai botnet data structures to identify the patterns in attacks, both from the source and destination perspective. According to Bhattacharya, the biggest challenge that she sees in the field of information security is the ability to predict attacks and recognize patterns of networks threats in order to keep the Internet as secure as possible.

CREATIVE

Shawn Doughty, Creative Direction
Brendan O'Hara, Art Direction & Graphic Design
Georgina Morales-Hampe & Billy Kamenides, Project Management



ABOUT AKAMAI

As the world's largest and most trusted cloud delivery platform, Akamai makes it easier for its customers to provide the best and most secure digital experiences on any device, anytime, anywhere. Akamai's massively distributed platform is unparalleled in scale with over 200,000 servers across 130 countries, giving customers superior performance and threat protection. Akamai's portfolio of web and mobile performance, cloud security, enterprise access, and video delivery solutions are supported by exceptional customer service and 24/7 monitoring. To learn why the top financial institutions, e-commerce leaders, media & entertainment providers, and government organizations trust Akamai please visit www.akamai.com, blogs.akamai.com, or [@Akamai](https://twitter.com/Akamai) on Twitter. You can find our global contact information at www.akamai.com/locations. Published 08/17