Promoting Stakeholder Action Against Botnets and Other Automated Threats

A Response Prepared for

DEPARTMENT OF COMMERCE

National Telecommunications and Information Administration
[Docket No. 170602536-7536-01]

RIN 0660-XC035

Akamai Technologies, Inc.

11111 Sunset Hills Road
Suite 250
Reston, VA 20190

Prepared By:

Micah Maryn, CISSP

Senior Solutions Engineer

703-581-6423

mimaryn@akamai.com

About Akamai

Operational since 1998, Akamai is the largest and most robust Content Delivery and Cloud Security Platform, providing enterprises across the globe secure, high-performing user experiences on any device, anywhere.

At the core of Akamai's solutions is the Akamai Intelligent Platform™, the largest, most distributed, FedRAMP accredited cloud infrastructure on the Internet. Consisting of more than 233,000 servers in over 130 countries and within more than 1,600 networks around the world, Akamai sits within one network hop of 85% of all client users (malicious and non- malicious actors), delivering over 30 terabits of traffic and over 3 trillion interactions daily. Akamai supports thousands of organizations, including:

- 60 percent of Global 500 & Fortune 500 companies
- The top 30 media & entertainment companies
- All 20 top global e-commerce sites
- 18 of the top 20 world's largest asset managers
- 8 of the top 10 world's largest FinTech firms
- Thirteen of the top 15 largest auto manufacturers
- Nine of the top 10 global pharmaceutical companies
- Six of the top seven computer manufacturers
- All of the top anti-virus companies
- All branches of the U.S. military
- Multiple Federal agencies within every cabinet level department

Akamai's government experience includes over 100 government agencies and a presence within each of the 15 cabinet level departments. In addition to our product solutions, several agencies have implemented customized which leverage our insight and intelligence capabilities.

Akamai extends the security perimeter to the edge providing unmatched reliability, security, and visibility. Using this far reaching visibility to monitor current Internet conditions and activities enables Akamai to aggregate real time threat intelligence, which is then applied to further strength the security of our customers and enhance our cloud security solutions. Real-time data feeds of our threat intelligence are also available and are currently being used by several federal agencies.

Using the reach of our platform, our extensive threat intelligence, and our practical experience with active mitigation, Akamai has been successful in mitigating hundreds of attempted DDoS and application-layer attacks. Akamai has mitigated at least one major attack per week for years. Major trends are discussed in our quarterly *State of the Internet Report* (the most recent report, Q1-2017, has been included as Attachment 1). Some notable events mitigated by Akamai include:

- September 11, 2001- flash mobs and loss of major Internet connections inside the World Trade Center crippled news sites, the sites delivered by Akamai continued to be operational.
- July 2009- Akamai defended numerous US Government, commerce, and financial services sites from a multi-day 124-Gbps DDoS coming from a bot-net within South Korea (likely attributed to North Korea) with no operational impact.
- December 2010- Akamai defended several of its customers against a highly-contested hacktivist campaign known as Operation Avenge Assange.
- September 2016- Akamai successfully mitigated one of the largest attacks ever seen, in excess of 665 Gbps targeting security blogger, Brian Krebs.[1]

[1] Many media outlets inaccurately reported that Akamai had dropped Krebs due to the impact on our platform. krebsonsecurity.com had been protected pro bono by our Prolexic scrubbing solution for 4 years

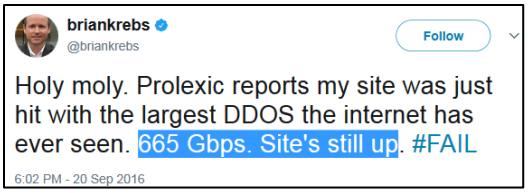

Twitter @briankrebs 9/20/2016

In addition to our core platform and threat intelligence capabilities, Akamai has assembled a team of security experts who are proactively engaged in increasing the security posture of our customers and our platform. Akamai operates five Security Operations Centers around the world maintaining 24x7 operations to support our customers during any security event. We also have teams dedicated to the analysis of the massive amounts of data aggregated to identify emerging threats and develop successful mitigations.

Akamai staff are certified across industry domains, such as CISCO networking certifications, Project Management certificates (e.g. PMP), and Security certifications (CISSP, GWAPT, GSLD, CEH, GPEN, etc.) to name a few. In addition, Akamai and Akamai employees actively participate with several working groups to solve problems like botnets, incident response, cloud security, and law enforcement (complete list in Appendix 2). Most recently, Akamai's CEO, Dr. Tom Leighton, and CSO, Andy Ellis, participated in the White House Technology Summit.[2]

Finally, Akamai has been a strong partner in support of information sharing both within the public and private sectors. Akamai is a member of Financial Services Information Sharing and Analysis Center (FS-ISAC), including providing cloud security services for the protection of the FS-ISAC web sites[3], and an active partner with the Department of Homeland Security's (DHS) National Cybersecurity and Communications Integration Center (NCCIC)[4]

(prior to and after Akamai acquired Prolexic). In that time, the Prolexic solution mitigated around 269 attacks (Source: Krebs 11/22/2016). We have included some links to articles which more accurately describe the events which led to that decision.

1. Motherboard.com interview with Brian Krebs
2. Boston Globe 9/23/2016

[2] NBC News: Tech Titans Meet at the White House

[3] https://www.fsisac.com/about/PoweredByAkamai

[4] https://www.dhs.gov/national-coordinating-center-communications

Cyber Security & the Dynamic Threat Landscape

Organizations today operate in a Faster Forward world. Over three billion people are connected to the Internet, often through multiple computing devices. People are spending a growing portion of their lives online – communicating, shopping, being entertained, and working. For both business and government organizations, this represents a significant shift in how they engage with their customers and employees.

For these organizations, more of their daily activities now take place outside of the traditional office. They engage with customers and collaborate with coworkers over the Internet, performing financial transactions, transmitting sensitive business data, and communicating over public networks. To do this, they are moving more of their applications onto Internet-facing networks, so customers can shop 24x7 and employees can access the resources they need at any time in the global work day.

There are two trends have made the threat landscape far more dynamic:

1. Mobile device usage
2. The rise of the Internet of Things (IoT)

The new challenge that mobile technology introduces is centered in the fact that a mobile device is, at its core, a computer that can operate inside and outside the secure network environment. It rises new concerns over the how to protect the data and information within a device and how to prevent the device from being used as a vehicle for introducing threats into the enterprise environment.

IoT now presents us with an extremely dynamic threat landscape. We now have multiple IP connected devices raging from cars to refrigerators. IoT provides great benefits for monitoring and management of a variety of devices and situations. At the same time, if left unsecured, IoT also provides an exponentially larger base which can be exploited by malicious actors to launch attacks. The two key security concerns with IoT are:

1. How to use IoT to your benefit, while also maintain the security of the IoT devices
2. How to protect your networks from outside IoT devices that have been compromised

APPROACHES TO CYBERSECURITY

The changing focus of security threats – from network to applications, disruption to data theft, and one-dimensional to multi-dimensional attacks – is driving an architectural shift in the security industry. While DDoS attacks will continue to command the greatest attention, many of the most damaging attacks are also the most difficult to detect, and provide little to no advance warning. This necessitates a security posture that is always on, but still provides the performance and scale to respond to the largest network- and application-layer attacks prevalent today.

ON-PREMISES HARDWARE

Many organizations rely on hardware devices, such as network firewall, DDoS mitigation, and web application firewall (WAF) appliances, deployed on-premises within their data centers. From a financial perspective, on-premises hardware requires a large upfront capital expenditure with a typical hardware lifecycle and depreciation of two to three years, as well as operational expenditures for IT management costs. Deploying applications across multiple data centers can further increase costs as these solutions often must be deployed wherever the applications are

located. While private sector organizations can be relatively quick in updating this infrastructure, the federal government continues to lag behind[5], increasing operational and maintenance costs, and increasing overall risk.

As with any inline solution, the challenge for on-premises hardware is ensuring sufficient scale and performance to remain resilient against attacks that are growing in size. This challenge is particularly acute for hardware devices, which are typically limited by the capabilities of the individual device, as opposed to those of the entire security system:

- **Scale** – While hardware devices are always increasing in scale, hardware-based security systems can still be overpowered by the vast amount of attack traffic generated by today's massive botnets. In order to defend an against 75% of current DDoS attacks, an organization organization's defenses would need to be able to mitigate 1.3 Gbps of volumetric DDoS attack traffic directed at its infrastructure. To withstand 95% of attacks, those defenses would need to be able to absorb an attack of 5 Gbps. It is important to understand that there are still a significant number DDoS attacks generating more than 100 Gbps of traffic. Such attacks are common enough to be a concern.[6]

  And it's not just bandwidth; the processing power of the devices is also a factor. While one may have an extremely large bandwidth capacity, the appliances themselves may not have a proceeding power to handle the millions of packets per second (MPPS).

- **Performance** – Defending against application-layer attacks can be extremely resource-intensive. Web application firewalls require a large amount of computing resources to compare incoming application traffic against known attack profiles. Even normal application traffic can require a significant amount of processing, which can reduce the published performance of hardware-based WAF devices and, subsequently, the amount of traffic that makes it through to the applications behind them.

When considering a hardware-based approach, it is important to remember each hardware appliance operates with a sequence chain of traffic flows and process. A DDoS attack needs only to overpower the weakest link in the chain in order to cause an outage for the entire system.

A final disadvantage of a hardware approach is that it attempts to stop a DDoS attack only after it has entered the data center. If an organization does not have a sufficiently large Internet link, then the attack will saturate the available bandwidth, causing an outage for the entire data center.

INTERNET SERVICE PROVIDERS

Another common approach to protecting Internet-facing applications is to implement a service through an organization's Internet service provider (ISP). In the federal government, many organizations are using specific ISPs certified as Trusted Internet Connections (TICs) as well as ISP which participated in the Einstein 3 program. Many ISPs offer DDoS mitigation services to complement their primary business of providing network bandwidth.

This approach can be attractive for several reasons. First, it enables them to transfer the responsibility for mitigating a DDoS attack to a third-party. Network traffic to the organization's applications is already transiting the ISP's network, making the ISP a logical location in which to implement a DDoS mitigation service. However, there are several considerations that may not be apparent at first glance:

[5] Government Accountability Office Report on Legacy Systems

[6] Akamai *State of the Internet Report Q1-2017* (Attachment 1), p.13

- **Multiple ISPs**–ISPs can only mitigate DDoS attack traffic transiting their network. Many government agencies are required purchase bandwidth from multiple ISPs in order to increase availability. But, this could significantly increase the complexity of responding to any DDoS attack. Attack traffic can now arrive over the networks of multiple ISPs, meaning that organizations must deploy and manage multiple DDoS mitigation solutions, as well as coordinate any DDoS attack response across multiple vendors.
- **Scale** – with the largest DDoS attacks exceeding 600 Gbps in peak bandwidth, most ISPs simply do not have sufficient network capacity to properly mitigate potential attacks directed at their customers. However, even smaller attacks can present a risk to the ISP's network, consuming network capacity and impacting performance for other customers. Faced with this situation, an ISP will often choose to "black hole" traffic to the intended target of any DDoS attack over 10 Gbps in size in order to preserve the stability of the ISP's network at the expense of the target organization.
- **Security Expertise** – most ISPs do not regard security as a core component of their business, but rather an additional capability to augment their primary business of providing network bandwidth. As a result, ISPs typically have limited security expertise and do not employ best-of-breed security solutions, and may have difficulty stopping attacks that are too complex or large in size.

Beyond DDoS mitigation, organizations must also consider the threat of posed by data exfiltration attempts such as SQL injections and XSS. Organizations that have deployed an ISP-based DDoS mitigation service must still implement a WAF solution in order to protect their websites and applications from data theft. This requires augmenting the ISP-based DDoS mitigation with a separate solution from another vendor – either an on-premises or a cloud-based WAF. In addition to the cost of acquiring and managing multiple solutions, this can increase the complexity of responding to multi-dimensional attacks that combine DDoS with data theft.

ENDPOINT SECURITY

Endpoint security tools and systems excellent for managing enterprise users' computers, ensuring that security policies are being enforced and scanning for potential threats with malware detection and antivirus software. However, malware and phishing attacks are become increasingly sophisticated especially when the threats are state sponsored. Some of these very sophisticated threats can go undetected for months or more.

To mitigate against these threats, organizations need to consider solutions which can complement endpoint security by looking at where requests are going, and using real time threat intelligence to prevent requests from potentially malicious destinations.

CLOUD BASED SECURITY APPROACH

Cloud-based security solutions provide a new approach to detecting and mitigating security threats. Here, organizations deploy a third-party cloud platform in front of their private infrastructure and inline between remote users and their websites and applications. The cloud security provider can examine network traffic for known attack patterns and pass only legitimate traffic through to the application. This extends the security perimeter beyond the traditional approaches, mitigating threats even before they hit the ISP. The further the reach of the cloud security platform, the larger the capacity, scale, and security perimeter.

For many organizations, the concept of stopping attacks in the cloud represents a paradigm shift. This approach moves the point of mitigation from the data center to the cloud platform and offloads a large part of the mitigation effort from an organization's IT staff to that of the cloud provider. This provides several advantages over traditional approaches:

- **Extending Security Perimeter** – defending against DDoS attacks within the data center requires scaling and hardening many infrastructure components. By moving the point of mitigation to a provisioned cloud platform, organizations can remove the complexity of securing every part of their infrastructure from different types of DDoS attacks.
- **Scale** – by leveraging the economies of scale that come from protecting many organizations at once, cloud providers can build a much larger infrastructure than what individual organizations can on their own. However, not all cloud security solutions are created equal – even between different cloud providers, the scale of their platforms can vary greatly. Organizations should evaluate the total capacity of the cloud platform – how much traffic it delivers on a daily basis as well as how much extra capacity it has to mitigate potential attacks and handle future growth.
- **Performance** – some cloud-based security solutions can improve performance while protecting applications against DDoS and web application attacks. These solutions often share a common underlying platform with a content delivery network (CDN) that is designed to accelerate access to web applications. Because many performance-sensitive applications may already be behind by a CDN, this approach can help secure those applications without requiring a tradeoff in performance.
- **Threat Intelligence** – cloud security providers typically have greater visibility into attacks and attack trends than individual organizations. Because of their position in the network, they can see an attack as it is first used against one of their customers and then leverage the technologies and techniques used to defend against that attack to improve security for other customers. Cloud security providers can make threat intelligence available to organizations in different ways, including through improved WAF rules, new attack signatures, customer-facing threat advisories, and better internal response processes.
- **Expertise** – the effectiveness of any organization's ability to respond to DDoS or web application attacks is greatly influenced by its experience at mitigating other similar attacks. By defending against attacks directed at many individual organizations over time, cloud security providers can develop significant expertise and experience. They can draw on this experience when mitigating future attacks to reduce mitigation times and any impact on their customers.
- **Compliance** – many organizations operate websites and applications that are subject to various legal regulations, such as the Payment Card Industry Data Security Standard (PCI DSS) for any site that handles credit card information. Within the Federal government, there are additional compliances which must be met, include FedRAMP and several OMB mandates. Organizations must ensure that their cloud security solution also complies with all applicable legal regulations to which they are subject.

PROTECTING ALL SIDES

External threats have three primary vectors for attack:

1. The Application- targeting the web application for data theft, defacement, or denial of service.
2. The DNS- targeting the DNS to prevent users and IP based services from being able to resolve the host.
3. The Core Infrastructure- targeting the host infrastructure

There's also the risk of internal threats. Common examples are phishing and malware which is inadvertently activated by users within an enterprise's network.

DENIAL-OF-SERVICE ATTACKS

One of the most common and pervasive security threats today is the denial-of service (DoS) attack. DoS attacks attempt to disrupt a critical Internet service by overwhelming a supporting infrastructure component, such as a web server or network device or consuming available network bandwidth. In order to mobile technologies to be an effective tool, the application server and hosting environment must be protected from lost availability from DDoS attacks.

Two trends are driving the increase in the size of volumetric DDoS attacks:

1. The growth in the traffic-generating capacity of large botnets, deriving from an increasing number of IoT devices and computing power of every connected device.
2. The continuous use of reflection and amplification attacks using vectors such as DNS, NTP and SSDP reflection. Reflection-style techniques exploit vulnerabilities in existing Internet services to generate much larger attacks than otherwise possible. For example, DNS reflection generates 28x to 54x amplification in attack size, while NTP reflection generates 556.9x amplification.

The rise of Mirai has made bot-nets even more dynamic. With Mirai, the bot-net operator has the ability to launch attacks over just about any port and protocol. Furthermore, Mirai has the ability to let the bot-net operator directly control the traffic generated rather than just create traffic through reflection and amplification.

This rapid growth highlights the difficulty in defending against volumetric DDoS attacks. Individual organizations can continue to invest in additional network bandwidth and higher performing network devices. However, they will always be hard-pressed to respond to the largest DDoS attacks of the time. These attacks harness the power of Internet to scale beyond the financial and technological resources of individual organizations.

DNS DDOS

DNS has become another favorite target for attackers, not just because of its critical role in the IT infrastructure, but also because it is typically one of its least scalable components. Many organizations only deploy a small number of DNS servers, making it vulnerable to a volumetric attack that could easily overwhelm it.

We have seen in the past how attackers have used amplification attacks to target an organization directly, or to use reflection to attack another. More recently we have seen Mirai use what is called *DNS Slow torture*, which sends requests for random subdomains.

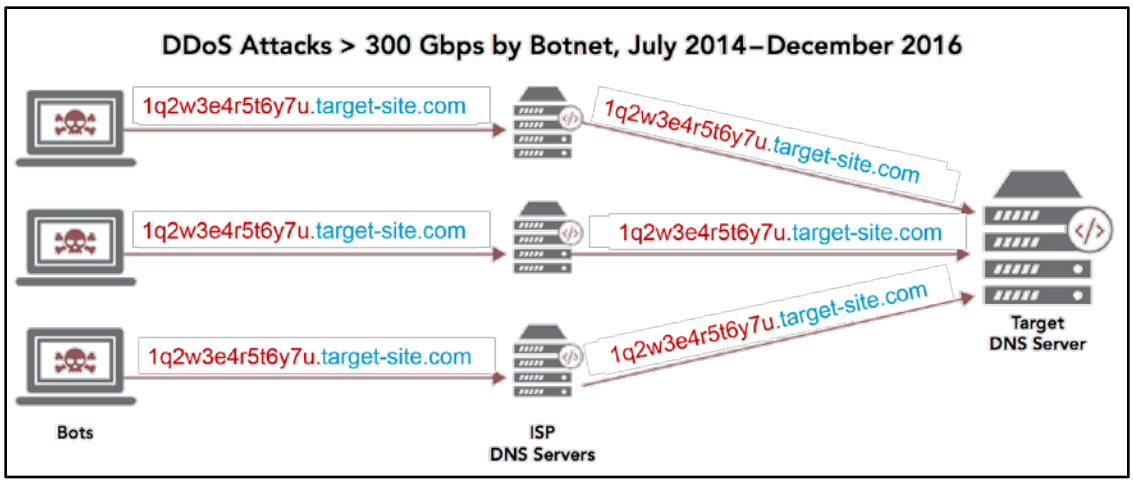

Mirai DNS Slow Torture

Since the query resolve would not have the non-existent, randomized subdomains in cache, each query is sent on to the authoritative resolver. While the DNS query flood generates relatively limited volumes of traffic, it still results in a denial of service outage by consuming the target domain's resources by forcing the look up of a high volume randomly generated domain names.

ATTACKS SHIFTING TO THE APPLICATION LAYER

DDoS attacks targeting the application layer may prove to be a more vexing long-term challenge. As advances in technology and user expectations grow so does the number of and complexity of web and other Internet-facing applications. Application-layer DDoS attacks come in different forms and use a variety of methods and techniques to deplete a web or application server of the resources it needs to operate. Two common examples of application-layer attacks include:

1. **DDOS –** As the underlying foundation for modern web applications, many application-layer attacks exploit HTTP vulnerabilities in order to incapacitate the targeted web server.
2. **Data breach-** breaching the security of the application to gain access to confidential information or even to deface the web site/application as a way to embarrass the targeted organization.

APPLICATION-LAYER DDOS ATTACKS

Application-layer DDoS attacks are more difficult to detect than network-layer attacks because traffic generated often looks legitimate. HTTP floods may generate high volumes of traffic, but to many traditional security tools, focusing on the network layer traffic, the HTTP requests appear legitimate.

In some cases, the attack is not even volumetric in nature. Rather, the attack is meant to exhaust the reroutes of the application. Attack like this might continue to open new connections, exhausting the number of concurrent connections that could be support. In other cases, the attacker could, take advantage of known bugs or vulnerabilities within an application by sending information or a command which will cause the application to fail.

TARGETING APPLICATIONS FOR DATA THEFT

Many web application attacks are designed not to disrupt operations but rather to steal data. As organizations today increasingly interact with their clients online, it is critical not ensure the integrity and confidentiality of the data.

Some common attack vectors for data theft include:

1. **SQL Injection –** According to Veracode, 30% of all data breaches are due to SQL injection.[7] This type of attack exploits web applications that do not properly sanitize user inputs and tricks them into running database code that returns more data than they otherwise would have.
2. **Remote File Inclusion (RFI)** – Similar to an SQL injection, this type of attack exploits web applications that do not properly sanitize user inputs. However, the immediate goal of a remote file inclusion is not to steal data, but rather trick the web server into executing the contents of a file stored in a remote location. In this manner, an attacker can then take control of a web server for malicious purposes.
3. **Credential Abuse –** Public-facing websites and applications often require users to log in to access parts or all of the application. Because users often use passwords that are easy to guess and share passwords across multiple accounts, hackers can purchase stolen user login credentials for one site and make repeated login attempts against other sites in order to compromise an account.

Web application attacks can be difficult to detect. SQL injections, remote file inclusions, and credential abuse attacks generate application traffic that appears legitimate to traditional network-layer security tools. As a result, organizations are often not aware of ongoing attacks until after large amounts of data have already been stolen. In an analysis of recent data breaches, Verizon found that 99 percent of attackers compromised their target within days or less, but only about 10 percent of breaches were discovered in that same time frame. [8]

A MULTI-DIMENSIONAL SECURITY THREAT

While many security solutions focus on defending against a single type of attack, attackers are increasingly employing multiple different types of attacks in combination. Multiple-dimensional attacks have a higher chance of succeeding against organizations that may have limited IT resources or solutions focused on a single category of security threats.

But even against well-protected applications, these attacks test their target's ability to respond to multiple parallel attacks occurring in different parts of their IT infrastructure.

In addition, attackers are beginning to combine DDoS attacks with SQL injections, using bandwidth-consuming and noisy DDoS attacks to distract limited security resources from the attacker's true goal of data or financial theft. This scenario highlights the danger of focusing on just one type of attack vector. In a rapidly changing threat landscape, organizations must be prepared to respond to a variety of potential attacks, including combinations of different types of attacks, in order to safeguard their IT infrastructure.

This is why leveraging provisioned cloud security services levels the playing field. By extending the security perimeter and providing dynamically scalable infrastructure to match the dynamic scalability of the new dynamic threat landscape, organizations would be able to mitigate these threats before they ever reach the application, the DNS, or host infrastructure.

THREAT INTELLIGENCE

---

[7] DuPaul, Neil (July 2013). The Real Cost of a Data Breach Infographic.
Retrieved from http://blog.veracode.com/2013/07/the-real-cost-of-a-data-breachinfographic/

[8] 2016 Data Breach Investigations Report. Verizon.
Retrieved from http://www.verizonenterprise.com/verizon-insights-lab/dbir/2016/

Another reason cloud based security is so effective is the ability to gather far more threat intelligence. When mitigating threats at the origin, one is only seeing the threats targeting the origin. However, new and emerging threats are always evolving; Attacks targeting someone else today, will likely target you tomorrow.

A highly distributed cloud security platform not only provides a scalable infrastructure for mitigating, but also a deeper view of current internet conditions, and emerging threats. The threat intelligence can help support predictive analysis, and complement solutions for the mitigation of internal threats.

MITIGATING THE INTERNAL THREATS

The other threat that needs to be address is the internal threats cause by phishing and malware attacks. While endpoint security tools are the primary way to address such threats, like bot-nets, these threats continue to evolve. Even with the most robust endpoint security, threats can still be introduced, and often by the inadvertent actions of the client within an enterprise.

MALWARE & PHISHING

Malware is large threat to organizations, in which attackers attempt to introduce software into the enterprise to disrupt computer operations, gather sensitive information, or gain access to private computer systems.

Malware can be introduced in various methods, such as a file inclusion on a website or software installer. Phishing has become a common method for inserting malware, especially to specific targets. In a Phishing attack, the attacker masquerades as a trustworthy entity in order to get the victim to inadvertently trigger the installation of malware.

The risks of Malware and Phishing become more compounded by mobile technologies and bring your Own Device (BYOD) policies. While the enterprise network maybe well protected, using malware detection and email filters to detect malware and phishing attempts, one still needs to protect the device form becoming compromised when the user is outside the enterprise network.

Users could fall victim to Phishing when using personal email, or they could install the latest fun new app, which includes multiple vulnerabilities or malware, which can now compromise the device, and potentially the enterprise network

CLOUD SECURITY & INTERNAL THREATS

Cloud security can benefits to the endpoint security in two ways:

1. Performing threat intelligence checks on outbound requests
2. Monitoring the flow of outbound requests for anomalous activities

Many malware infections are introduced by user's inadvertently click on a link or a file which then makes an outbound call to pull in the actual malware. Cloud Security solutions using a strong threat intelligence repository can be used to verify where the request is being directed. If the requests is resolving to an active threat, or even somewhere that is outside the enterprises acceptable use polices, then that requests would be blocked, preventing the malware from being pulled into the network.

Using that same approach, the cloud security solution can log and detect anomalous connections flowing out of the enterprise which may indicate an active malware threat.

Akamai Responses to RFC Questions

Respondents are invited to respond to some or all of the questions below:

1. *What works: What approaches (e.g., laws, policies, standards, practices, technologies) work well for dealing with automated and distributed threats today? What mechanisms for cooperation with other organizations, either before or during an event, are already occurring?*

   The increased usage of Cloud Service Providers (CSP) has been an effective approach for cyber security. In terms of threat mitigation, in particular threats caused by highly distributed bot-nets, leveraging a CSP provides organizations the additional capacity to mitigate such threats. In addition, cloud services help to ensure availability by providing automated disaster recovery and continuity of operations services. Essentially, cloud services are able to provide and additional layer of cyber security outside of an organizations infrastructure or ISP. How far that layer extends depends on the particular CSP, the specific services being used, and the architecture of the cloud service provider's infrastructure.

   A good example of the security benefits provided through CSPs is Akamai's support of the U.S. State Department in the implementation of a cloud-based web security solution. Since implementation, the solution has been successfully denying approximately 10 million malicious connections per month and offered full protection against one of the largest distributed denial of service (DDoS) attacks against the agency.[9]

   For federal organizations, which have required baseline security controls, procuring cloud services had been a cumbersome task of due diligence. FedRAMP has been en effective program for simplifying due diligence for procurement of Cloud Services for federal organizations, by accrediting multiple CSPs as meeting specific baseline security standards. This has resulting in federal agencies being able to quickly determine qualified CSP candidates, reducing the time and cost of procurement cycles.

   DHS's National Cybersecurity and Communications Integration Center (NCCIC) has been effective in aggregating threat intelligence and disseminating information about active threats and vulnerabilities. Acting as the primary federally centralized organization for cyber security, it is now effectively collaborating with 64 private and 11 federal agencies.

2. *Gaps: What are the gaps in the existing approaches to dealing with automated and distributed threats? What no longer works? What are the impediments to closing those gaps? What are the obstacles to collaboration across the ecosystems?*

   When dealing with automated and distributed threats, security solutions and mitigations need to be able to scale rapidly to meet that threat. In order to be scalable to match distributed threats, the mitigation platform needs to be distributed.

[9] Technology CEO Council (TCC) Government Efficiency Report, January 2017, p. 16

It is no longer possible to fortify the data center, or rely on the ISP to mitigate the size of attacks that can be generated by highly distributed bot-nets, like Mirai. Attacks of this scale and magnitude need to be mitigated far upstream from the target and as close to the sources as possible.

That is not to say the origin side mitigation, or ISP mitigations are obsolete; rather it has become necessary to extend security beyond those boundaries. It is simply not economically feasible to continue expanding the infrastructure necessary to manage the size of attacks we have already seen from attacks like Mirai.

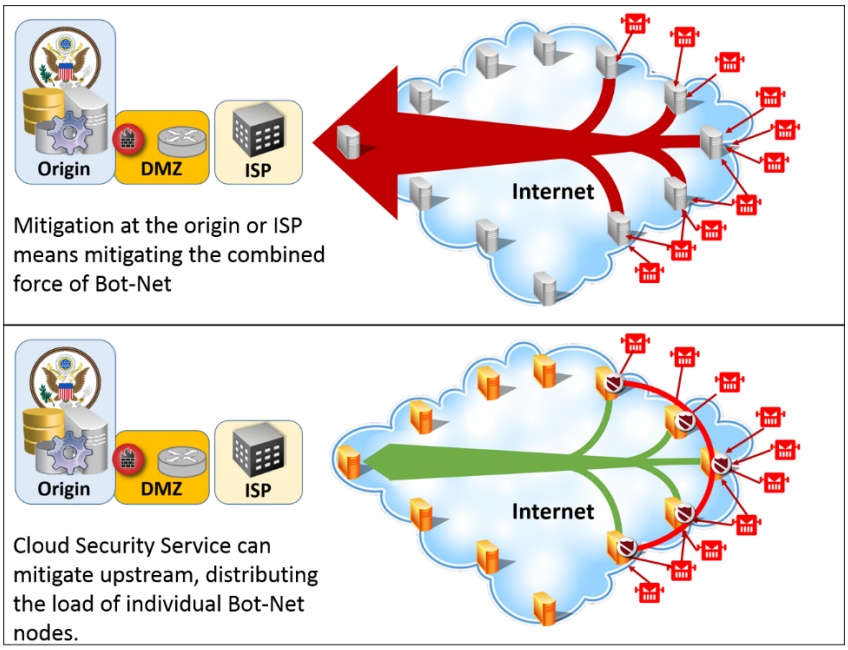

In the private sector, there a few impediments to procuring a cloud based mitigation service. Much of this has been a result of the competitive needs to reduce operational costs and reduce inefficiencies. This creates a need to invest in new technologies in order to reduce operational costs.

Private sector organizations spend roughly 15% on operations. By comparison, the federal government is spending about 30% on operations in support of mission delivery.[10] And much of that cost goes to support operations and maintenance (O&M) of existing systems. About 75% of the federal IT budget is supporting legacy systems, some of which are 25 years or older.[11]

According to the Government Accountability Office (GAO) 75% of all federal IT investment has been spent on O&M. As a result of the high cost of O&M, federal investment in IT development, modernization, and enhancement activities, has declined by $7.3 billion since 2010.[12]

[10] Technology CEO Council (TCC) Government Efficiency Report, January 2017, p. ii

[11] Nextgov.com February 9, 2016 *"White House Wants to Give Agencies New Pot of Money to Upgrade Legacy IT"*

[12] Government Accountability Office Report on Legacy Systems, May 2016, p. 2

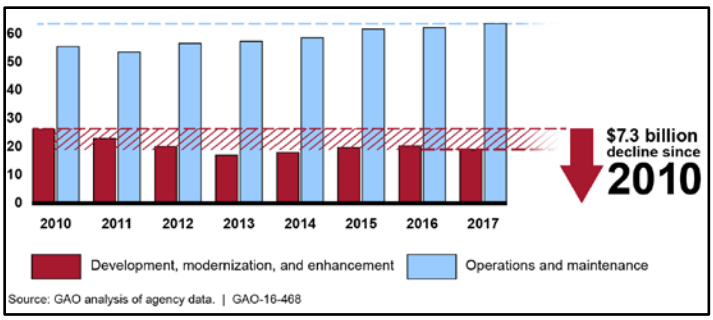

The GAO report breaks down the numbers further, identifying *Provisioned IT Services* such as cloud services and shared services, and *Non-Provisioned Services*, developed and maintained within an organization.

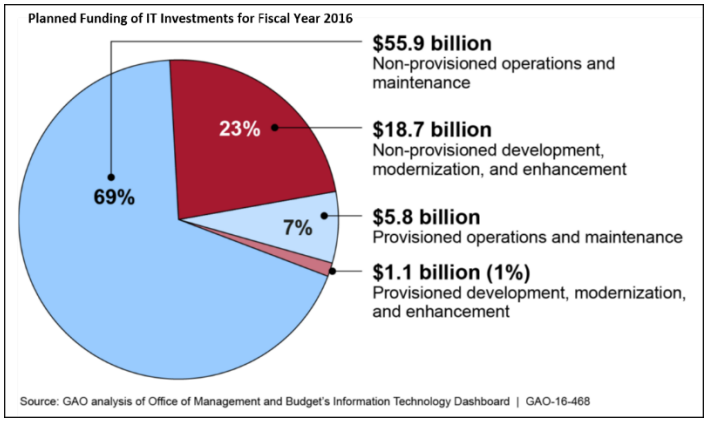

The GAO analysis indicates that only 24% of FY 2016 spending had been planned for investment into development, modernization, and enhancements (DME). Only 8% had been planned on investment into provisioned services- only 1% for DME of provisioned services.[13]

In order to match the size and scalability of bot-net threats, the government must increase IT investment into Provisioned IT services. DME investment into non-provisioned infrastructure would best be spent in conjunction with DME investment into provisioned IT services. Specifically, cloud security services.

Investment now in provisioned cloud security services will save more in future O&M and the resulting cost caused by security events. The average cost of a DDoS attack is $40,000 per hour with an average cost per incident around $500,000.[14] The post event cost of Office of Personnel and Management (OPM) data breach has been estimated as being $350 million or more.[15] By investing now to prevent these attacks, will significantly reduce the cost of security events; money which could then be spent on additional DME.

Investment into provisioned cloud security services often times has additional cost savings benefits. The obvious savings is a reduction in the size of non-provisioned infrastructure, and the O&M cost for that infrastructure. Furthermore, provisioned services would significantly reduce

[13] Government Accountability Office Report on Legacy Systems, May 2016, pp. 4-5

[14] Securityweek.com article

[15] Fedtech.com article

the future costs for upgrade non-provisioned infrastructure. In addition to reducing infrastructure costs, provisioned cloud security services may also have the added befit of improving performance and the user experience, both of which will provided cost savings in other areas.

For example, Akamai's cloud security solutions can be provisioned with features for optimizing the performance and optimizing the user experience by adapting content to the client user's device. According to a 2015 Pew research study, 7% of American households are *Smartphone Dependent,* meaning their only way to access the internet is via a smartphone or tablet. Furthermore, the study found that 40% of all American use a smartphone or tablet for accessing government information and services.[16] A 2016 survey by CFI group found that users of government low cost mobile applications had a higher customer satisfaction rates than those who had used more costly call center support.[17] Provisioned cloud security services which can offer additional performance and optimization features would not only improve security, but also reduce costs for the DME and O&M for mobile adaptive technologies as well as reduce O&M for more costly call centers.

An additional obstacle to government acquisition of provisioned cloud based security has been the government procurement process. The technical team of an organization soliciting cloud security solutions needs to go through an agency or departmental Contracting Officer. During the initial phases the technical teams will create a set of requirements, sometimes talking with various vendors. Once the requirements have been finalized, either a solicitation for proposal will be issues, or, solicitations may be sent out to several vendors for quotation.

Once a quote has been provided, very often negotiations are managed by the Contracting Officer. Unfortunately, contracting officers are not subject matter experts with respect to the technical buyer's needs. The result is the solutions are reduced in their effectiveness in order to reduce costs.

Another issue is an apparent restriction for procurement based solely within the scope of a solicitation. Often vendors are able to identify additional solutions to further enhance the efficiency of the system described within a solicitation or perhaps identify security gaps within a solicitations technical requirements. The vendor may describe the gap within their proposal and including an additional solution to resolve that gap. Or, going back to our previous example regarding performance optimizations, a vendor may propose a provisioned cloud security solution which also helps reduce the O&M cost for web mobile or mobile app support. Unfortunately, since these additional features are outside of the immediate scope of a solicitation, contracting officers and are restricted from (or unwilling to) authorizing the additional features, regardless of the potential increase in efficiency or reduction in on going O&M.

Our recommendation would be to enable the technical buyer's to have more influence in these situations so they are able to procure the solutions which fully meet their security requirements. Second, enable procurement to evaluate the potential value of features which have clear cost reduction benefits or improvements, but are outside the scope of the initial solicitation.

---

[16] Pew Research center, U.S. Smartphone Use in 2015

[17] CFI Group Government Contact Center Satisfaction Index (GCCSI) 2016

Trusted Internet Connection (TIC) requirements and reporting requirements of programs like Einstein 3 (E^3A) have been an obstacle.

Some organizations view the reporting requirements of the TIC and E^3A programs so strictly, it prevents them from using provisioned cloud solutions like hosting, CDN, and, most important for the purpose of this RFC, Provisioned Cloud Security solutions. The concern being that by leveraging a provisioned cloud security solution to mitigate threats in front of the TIC would result in a failure to record security event data, thus preventing the aggregation of any actionable intelligence by the TIC. As a result, these organizations are relying on their own infrastructure and the TIC/ E^3A provider (i.e. the ISP) to have the scalability to manage a high volume event.

There are also concerns with how efficiently E^3A provides mitigation for protected threats. The E^3A MTIPS provider is capturing inbound data and sending it off to DHS for analysis. After analysis, DHS pushes out new rules for what types of signatures need to be flagged or blocked. However, such analysis would take time, and would not necessarily provide mitigation for new threats which are actively taking place.

Cloud Security Solutions have proven to be an effective and valuable tool for increasing the scalability and mitigating the threats posed by highly distributed bot-nets. It would make sense to enable organizations to leverage Cloud Security Solutions without losing the ability to meet the reporting requirements of the TIC or E^3A programs.

Our recommendation would be to create a standardized policy and process for providing the data captured by upstream cloud security solutions to DHS. By having a standardize process for reporting on upstream mitigations would enable organizations to leverage provisioned cloud security services without concerns regarding reporting requirements. Furthermore, this approach will enhance the amount of threat intelligence being gathered by DHS.

Reducing the attack surface of the federal government by reducing the number of Internet Access points is another place where the government could improve. A first step would be to logically divide government networks into three categories: government user access networks; government application networks, and government information networks.

1. Government user access networks should be limited to the devices used by government employees, and the immediate infrastructure (email, etc) that supports those devices. These networks should be limited in what traffic is permitted inbound, and in lateral movement - consider these as a "secure ISP. "
2. Government application networks are for application services that are available to the government users. While access might be limited to traffic coming from government user access networks, strong authentication - based on public-key-cryptography and out-of-band push - should be used to verify users.
3. Government information networks, designed to be accessed by the public, should be entirely distinct from the government application networks (Examples: DNS, inbound email, Pl websites). These should be public-facing, and designed to be hardened in the same fashion as e-commerce and financial services websites, and using commercial cloud systems where appropriate. Where information needs to flow between systems in these networks and systems in government application networks, the connections should be tightly controlled and mutually authenticated.

Organizations using these categorizations will help them better identify the applicable standards

and best practices necessary for the security of those systems. In turn, agencies would be able to make more informed procurement decisions.

Along with limiting government Internet access points, is limiting the number of inbound VPN connection. VPN is a great way to secure and grant remote access for federal employees and contractors. But it also provides the users access to the enterprise network, when really all that is needed is access to specific systems and applications. The system is only as secure as the remote systems granted VPN access.

The Target breach is the best example of the risks of VPN. Target was not directly breached. Fazio Mechanical, a small heating and air conditioning firm contracted by Target was breached. It was from the breach of Fazio Mechanical, that the attackers were able to steal the virtual private network credentials that Fazio's technicians used to remotely connect to Target's network.[18]

With VPN access being granted to multiple contractors, there's an increased risk to lateral movement within the network. Fort the most part, IT security can only monitor VPN login activity, not network activity.

Cloud security solutions can address this as well, but providing solutions which limit remote users to only specific systems and applications for which they have privileges.

Akamai offers such a solution today with user's access applications through Akamai's cloud platform, which secures user access far outside your network with no direct path into the network. Rather, Akamai's Enterprise Application Access solution is connected by a secure, outbound, mutually authenticated TLS connection initiated from within the enterprise network (or cloud host environment) and brings the application to the user's browser.

Since there are no tunnels, there is no path for malware to land inside your network and potentially spread to sensitive or privileged systems. All user connections are stopped in the cloud, terminating on secure proxies while applying strong authentication and security controls. You can add your own security controls for increased protection of highly sensitive applications.

In terms of enhancing collaboration efforts, the NCCIC should remain the primary federal centralized organization for the collaboration of threat intelligence. However, more does need to be done to support such efforts. The GAO released a report in February[19] outlining several key recommendations for improving the effectiveness of the NCCIC and the ability collaborate with other organizations. Some of the recommendations would be resolved by changes in policy and procedures. However, some recommendations, such as ensuring the integration of multiple data sets and ensuring the accessibility of information maybe more quickly adopted through collaboration of private sector organizations and systems integrators.

3. *Addressing the problem: What laws, policies, standards, practices, technologies, and other investments will have a tangible impact on reducing risks and harms of botnets? What tangible steps to reduce risks and harms of botnets can be taken in the near term? What emerging or long term approaches may be promising with more attention, research, and investment? What*

[18] https://krebsonsecurity.com/2015/09/inside-target-corp-days-after-2013-breach/

[19] http://www.gao.gov/assets/690/682435.pdf

are the public policy implications of the various approaches? How might these be managed, balanced, or minimized?

Since getting consensus on new legislation will be difficult and time consuming, efforts to mitigate the risks of bot-nets should first focus on efforts which would not require new legislation. That said, the current Modernizing Government Technology Act of 2017[20] (MGT Act), which passed the House and is now before the Senate, should encourage the adoption of cloud services and improvements necessary for cybersecurity.

We cannot stress enough, that *Distributed Threats* require *Distributed Mitigations*. Whenever applicable, organizations should invest in cloud security solutions with the capabilities of mitigating attacks as far upstream from the organizations infrastructure and as close as possible to the source. These solutions will provide increased and dynamic scalability far beyond what is possible with origin side or ISP investment.

This is a proven approach which Akamai is currently providing to thousands of customers, including hundreds of federal agencies.

Several things can be done to help ensure overall security and increase adoption of provisioned cloud services.

- Mandating widespread adoption of strong default, ubiquitous encryption for all government communications
- Improving identity management, including requiring multi-factor authentication for access to all networks with a near-term goal of eliminating password-based authentication entirely
- Protecting federal data through role and policy-based access controls

It is important that the government evaluates all commercial cloud providers to ensure that the cloud providers are able provide solutions which meet the demands of the federal government. This includes:

- High Availability & Scalability: The Cloud Provide must be able to scale as to ensure the availability during both DDoS attacks and unanticipated flash crowds
- Proven defense in depth security offering
- Compete end to end business owner visibility
- Full accreditation and compliance with base line security standards
    - FedRAMP in all cases
    - Any additional compliances which are applicable to the organizations mission

Not every CSP meets accreditation and compliance standards, or they offer some services that are accredited or compliant, but the actual service being used by the government is not operating within the accredited/complaint infrastructure.

Furthermore, it is critical for any organization looking at using provision cloud services to review the Accept Use Policies (AUP) of their potential CSP. Several CSPs have very permissive AUPs which allow their services to be used by hacktivists, malware distributors, and even terrorist organizations. Moreover, some CSPs are quite defensive about hosting and protecting such customers. This risk of using a CSP willing to support malicious actors goes beyond public embarrassment, it is also a risk to the availability of all sites, application, and content using the same platform.

[20] https://www.congress.gov/bill/115th-congress/house-bill/2227

The no-ip.com case is the best example of the risks posed by using a CSP which supports malicious activity. In 2014, The US District Court of Nevada ruled that Microsoft could seize a low cost web hosting platform, NOIP.com. Microsoft argued that NOIP.com was allowing cyber-criminal to use the NOIP.com platform to launch attacks and Malware at Microsoft and Microsoft's customers.

The court agreed with Microsoft, and allowed Microsoft to seize the NOIP.com platform by taking over NOIP.com's DNS. While this did stop the criminal activity, it also took down thousands of legitimate users of NOIP.com[21].

This case set a serious precedent, especially because it was a civil case resulting in a private company being allowed to take over an entire cloud platform because of the activities of customers using the same platform.

Any cloud services vendor willing to host illegal content or allow criminal activity on their cloud platform increases the risk to all the users on the same platform. Legal actions against the CSP could result in losing availability of the sites and applications, access to any content stored within the CSP, even losing the enterprise's DNS resolvers.

Adopting standards for the use and deployment of IoT devices is critical to improving the threat landscape in the long term. But today there are few market incentives to create devices with strong security; like any new technology, security/safety is typically one generation behind.

Akamai has been actively engaged with solving some of these IoT security concerns with auto manufactures. Specifically, in the areas around global content distribution and versioning control and the ability to allow auto manufacturers to perform OTA (over the air) updates for their fleet cars and consumer vehicles by leveraging Akamai's geographically diverse presence of its CDN system to deploy. In the process of addresses some of the key performance requirements necessary for the deployment of 1GB updates to all consumer cars, in the timespan of 48 hour, Akamai has been developing authentication methods to ensure that vehicles are only updated by the manufactures. Moreover, enduring the protection of consumer PII.

Because IoT is rapidly evolving, research needs to be done before a full set of standards could be published. Within the federal government, NIST should lead this effort in determining what those standards need to be.

But, simple things can be done now. A large part of preventing the threat posed by IoT bot-nets like Mirai, is preventing IoT devices from being compromised. When Mirai first emerged, it compromised many commercially available devices by simply using the manufacture's default authentication credentials. Following some basic guideline can help mitigate such risks, including:

- Always change factory-default credentials of any Internet-connected device
- Unless required for normal operation, completely disable the SSH service on any Internet connected device. If SSH is required, put "AllowTcpForwarding No" into sshd_config.
- Consider establishing inbound firewall rules preventing SSH access to your IOT devices from outside of a narrowly trusted IP space, such as your own internal network.

[21] http://krebsonsecurity.com/2014/07/microsoft-darkens-4mm-sites-in-malware-fight/

- Consider establishing outbound firewall rules in place for IOT devices at your network boundary, preventing tunnels established from resulting in successful outbound connections.

Education of consumers regarding best practices of password management and having organizations in both public and private sectors enforce standard username/password policies on newly deployed IoT devices is a good start as more comprehensive standards are established.

Ultimately, any legislation, mandate, standard, or policy will only be effective if there is some accountability for failure to meet compliance. A common argument against compliance is that the legislation, mandate, standard, or policy is "unfunded". While this is an understandable concern, CIOs do have options in meeting the requirements of unfunded compliances. Both the GAO and OMB have put forward recommendations for moving more infrastructure and services to provisioned and shared services offerings. Both have concluded that this would decrease the ongoing cost for O&M of non-provision infrastructure. Then take into account that that 69% of federal IT budgets are used for O&M of non-provision systems, it is reasonable to conclude that moving systems to provision cloud infrastructure should reduce the cost of O&M and enable federal organizations to meet these unfunded compliances.

4. *Governance and collaboration: What stakeholders should be involved in developing and executing policies, standards, practices, and technologies? What roles should they play? How can stakeholders collaborate across roles and sectors, and what should this collaboration look like, in practical terms?*

NIST

NIST will have a critical role in defining and setting standards for which any new legislation, mandate, standard, or policy would be based.

OMB

Must tie IT priorities to agencies' budget requests and allow CIOs to exercise discretion regarding agency level IT investments to create more cloud adoption.

GSA

GSA will continue to have a central role in procurement. This is especially true with respect to cloud service providers because GSA manages the FedRAMP program. Should the MGT Act be passed by the Senate, and ultimately signed into law GSA will have a larger role. The current version of the bill in front of the Senate, "…establishes a Technology Modernization Fund for technology related activities, to improve information technology, and to enhance cybersecurity across the federal government. The fund shall be administered by the Commissioner of the Technology Transformation Service of the General Services Administration in accordance with guidance issued by the Office of Management and Budget." [22]

OMB, GSA and NIST

Needs to issue guidance to federal agencies on best practices used to transition to cloud services and migrate legacy systems.

DHS

[22] Statement from MGT Act Summary

NCCIC is already established, so it would make sense to continue with improving NCCIC rather than institute a new center for collaboration. The GAO report previously cited[23] has made some good recommendations for consolidating the methods for reporting and aggregation, as well as address accessibility concerns for international partners.

In addition to the recommendations of the GAO report, the government should consider having NCCIC be able to better serve specific needs of certain agencies. For example, The Department of Health and Human Services announcement that it will establish a cybersecurity collaboration and education center for the health care industry[24]. The concern here is that by creating a separately managed system for cybersecurity management will cause confusion among both public and private sector organizations with respect to reporting and gathering reports. Furthermore, it would create yet another system for reporting and aggregation- one of the key issues raised by the GAO report.

Rather, the government should look at what DHS NCCIC could do to support the concerns of HHS and maintain an effective center for cybersecurity collaboration and communication.

5. *Policy and the role of government: What specific roles should the Federal government play? What incentives or other public policies can drive change?*

The key role of the federal would be in defining acceptable standards, policies, and compliances which are relevant to the federal government. It also means that agencies like GSA and OMB are taking steps to ensure that agencies are in compliance.

There needs to be accountability. Often Federal IT organizations are concerned with having their budgets reduced should they find ways to reduce costs. The result being more spending on O&M legacy systems than DME on more secure and cost effective solutions. This notion of "used it or lose it" needs to be changed. Rather there needs to be incentives for organizations to migrate off legacy systems to more secure systems.

For example, if an organization successfully migrates from a legacy non-provisioned system to a more secure provision system, they should be empowered to continue their efforts by reallocation of budget previously reserved for O&M of the legacy provisioned system to DME into making the new system even more effective.

Another approach maybe to reduce budget allocations in support of non-provisioned legacy systems. This approach would require agencies to find more secure and cost effective solutions.

6. *International: How does the inherently global nature of the Internet and the digital supply chain affect how we should approach this problem? How can solutions explicitly address the international aspects of this issue?*

Given the global nature of technology manufacturing it is important to understand what network devices should and should not be doing. This means ensuring that network traffic within the enterprise network and going out from the network enterprise secured and authorized. Our previous recommendations apply just as much to devices, including

- Mandating widespread adoption of strong default, ubiquitous encryption.

[23] http://www.gao.gov/assets/690/682435.pdf

[24] http://fcw.com/articles/2017/06/21/hhs-cyber-center-hearing-hsgac.aspx

- Protecting federal data through role and policy-based access controls and ensuring autonomous systems and devices are using mutual authentication
- Encouraging “security by design” through systems analysis methodologies like STPA-SEC

This also means network traffic within the enterprise network and going out from the network enterprise is being monitored for unauthorized connections. Within the enterprise this typically falls within the scope of endpoint and network security. For outbound connections, cloud security could be applied to perform a real-time threat intelligence check on outbound communications before the requested connections are resolved.

Using both approaches, agencies will be better able to detect potential compromises due to malware or other vulnerabilities.

We previously outlined things that end users can do to better secure IoT devices. But there’s also thing that manufactures should be doing. Unfortunately, we cannot always rely on the manufactures to improve security. Until firm IoT standards are created, it’s recommended that we ensure that IoT devices which are integrated within the network:

Device vendors:

- Do not have undocumented accounts set by the manufacture or vendor
- Have unnecessary ports and protocols disabled by default. For example, disable SSH on devices unless absolutely required for normal operations
- The ability to configure SSH to disallow TCP Forwarding
- Have a secure process for end-users to update sshd configuration so that they may mitigate future vulnerabilities without having to wait for a firmware patch.
- Ideally, Force users to change factory default account credentials after initial installation

7. *Users: What can be done to educate and empower users and decision-makers, including enterprises and end consumers?*

With the creation of the American Technology Council, the importance of training and education is key for making sure both consumers and enterprises have consistent messaging information.

There are guidelines and training resources available which can be shared with users to assist in adoption of best practices, and general education. Some of the resources are established from organizations such as SANS , OWASP, and NIST, which all publish regularly updated material on Information Security awareness, architectural guidelines, home user education, and online safety best practices.

With the rise of IoT in the home, it is important to educate the home consumer with information regarding very basic best practices. Probably the most important being how to secure one’s home IoT devices with strong user name and password. As stated previously, many IoT bot-nets are have compromised consumer IoT devices by simply using manufacture defaults to gain access. Simply setting strong username and passwords on consumer IoT devices will help mitigate the potential number of infected IoT devices, as well as protect the consumer from any potential data theft.

Appendix 1: Akamai's Cloud Security Approach & Solutions

Akamai offers an inline cloud security solution based on our Akamai Intelligent Platform. Originally founded as the leading CDN, the Akamai Intelligent Platform has evolved beyond acceleration to provide network- and application-layer security for websites and other Internet-facing applications. Its global scale and connectivity provides several inherent advantages when defending against many of today's most prevalent security threats.

AKAMAI INTELLIGENT PLATFORM

Akamai's Intelligent Platform consists of approximately 233,000 servers deployed in over 130 countries on 1,600 networks. Over 70,000 of those servers are within the United States. The Akamai's Intelligent Platform delivers 15-30% of the world's Web traffic daily. At the core of all of Akamai's performance, security, and availability solutions is the common Akamai Intelligent Platform.

Unlike other Content Delivery Networks which use a centralized architecture or "Super Pop" architecture, which route traffic thousands of miles away from the end users, Akamai uses a Distributed Architecture which brings content within one network hop of the end user. In fact, 85% of end users globally are within one network hop of the Akamai edge.

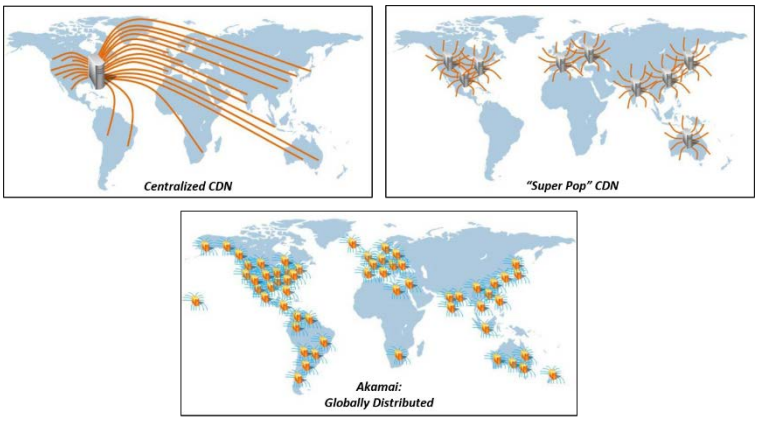

A Natural Architecture for Web Security

As a cloud-based security solution, the Akamai Intelligent Platform sits in front of websites and other Internet-facing applications, delivering network and application traffic from users to applications, and content from applications back to users. Its inline and distributed architecture provides two advantages when defending against both network- and application-layer attacks:

1. Inline – The inline architecture offers a natural location from which to defend against any type of DDoS or web application attack. As traffic passes through the Akamai Intelligent Platform to the application, the platform can identify and analyze attacks as well as take the appropriate actions to mitigate them. In addition, its inline architecture enables the Akamai Intelligent Platform to apply both positive and negative security models as appropriate for additional flexibility.
2. Distributed – Users access websites and other Internet-facing applications through the Akamai Intelligent Platform's globally distributed resources, including over 200,000 edge servers and seven global scrubbing centers. This provides a distributed platform for securing Internet-facing applications, with many locations in the network where mitigation activities can be performed.

Multiple Perimeters of Defense

Websites and other Internet-facing applications depend on a variety of infrastructure elements in order to function. These include the physical servers on which they run, the network infrastructure through which they communicate, and even the DNS infrastructure that directs client systems to the application. Protecting applications from downtime and data theft requires protecting all of these supporting elements from potential attack – a task that has become increasingly challenging as the IT landscape has shifted. Globalization and the resulting distribution of IT assets around the world, the adoption of cloud services and infrastructure, and increasing reliance on the Internet for business operations have all contributed to a diffusion of the traditional IT perimeter.

Akamai architected the Akamai Intelligent Platform as a distributed cloud platform in order to help organizations better protect their new, smaller, and more diffused perimeters wherever their IT assets are deployed and data is stored. The Akamai Intelligent Platform comprises multiple different technologies and networks that protect different parts of the application infrastructure, including:

- **Websites and Applications** – with over 233,000 servers deployed in 130 countries and over 1,600 networks, Akamai's edge network extends from the website or application to within one network hop from 85 percent of all client users. This provides Akamai with global reach to detect and stop both DDoS and web application attacks at the edge of the network, closest to where they begin and before they reach their target.
- **Origin Infrastructure and Non-Web Applications** – with seven high-capacity scrubbing centers located within key continental Internet hubs around the world, the purpose-built Prolexic DDoS mitigation network provides the capability to protect the entire origin infrastructure from DDoS attack. It employs over 20 different security technologies to detect, identify, and mitigate any type of DDoS attack targeting both the infrastructure as well as any type of Internet-facing application.
- **DNS** – an independent DNS platform architected for both performance and availability, Akamai's DNS platform includes thousands of name servers deployed in over 200 points of presence around the world to improve DNS performance and provide the capacity to absorb the largest DNS-based DDoS attacks.

Internet Security with Global Scale

Akamai architected every aspect of the Akamai Intelligent Platform for a hyperconnected world, with the capacity to handle network traffic on a global scale:

- On any given day, the Akamai edge network delivers between 15 and 30 percent of global web traffic. On any given day, the Akamai edge network delivers around 30 Tbps of web traffic, with far more available capacity. This Internet scale provides a natural advantage when defending against the largest DDoS attacks.
- The Prolexic scrubbing network provides over 3.2 Tbps of network capacity dedicated to mitigating DDoS attacks; almost 5 times the size of the attack targeting krebsonsecurity.com- 665 Gbps.
- The typical amount of traffic on Akamai's DNS platform represents less than one percent of its overall capacity, with spare capacity to absorb the largest DDoS attacks, including the 90 Gbps attack against a media company.

Beyond bandwidth-intensive DDoS attacks, the scale of the Akamai Intelligent Platform also provides a better defense against web application attacks. Detecting these attacks requires significant processing

power, as every incoming application request must be compared to known attack profiles through matching rules on a WAF. With over 233,000 servers distributed around the world, Akamai's cloud platform has the capability to protect against application-layer attacks without degrading the performance of the web applications behind it.

By leveraging the Akamai Intelligent Platform, organizations no longer need to plan to defend against the largest potential attacks. This allows them to reduce their capital and operational expenditures for on-premises hardware and network bandwidth. And when attacks do occur, the Akamai Intelligent Platform mitigates the attack at the appropriate network location in the cloud before it reaches the application, helping organizations maintain the availability and performance of their Internet-facing applications for legitimate users.

Always-on Security

Originally designed to deliver network traffic on a global scale, the Akamai Intelligent Platform provides a notable advantage over other security solutions – it is always on. Many solutions provide a passive and reactive defense. The target organization must first detect an attack before it can contact the security vendor to enable DDoS protection. Not only does a window exist in which applications are impacted, but this type of solution cannot effectively protect against many application-layer attacks that focus on data theft and blend in with legitimate traffic to go undetected.

The Akamai Intelligent Platform already delivers between 15 and 30 percent of all web traffic on a daily basis. It can inspect incoming network traffic for attack profiles while delivering it to the web application, providing both acceleration and security. With Akamai, IT organizations do not need to know that they are being attacked before they can defend against them. Akamai provides proactive Internet security that automatically detects new attacks as they begin, before they impact the target application, and without any outside intervention.

Protect & Perform

Most security solutions were designed for a single purpose – to defend against one or more types of attack. Because of this narrow focus, these solutions require organizations to tradeoff performance for security, resulting in lower traffic, lost lead conversion, and potentially reduced brand equity. For example, deploying hardware-based WAF can result in significant performance degradation for web applications. As a result, organizations often choose to deploy these security solutions out of band, despite the original design and greater security benefit of an inline solution.

Unlike many security solutions, the Akamai Intelligent Platform is architected with both security and performance in mind. Akamai views security and performance as complementary goals and helps organizations both protect and perform – protect web applications without requiring a tradeoff in application performance. The wide breadth of acceleration technologies also available for the Akamai Intelligent Platform allows it to protect web application infrastructures while improving application performance in order to maximize revenue and productivity at all times.

IMPROVING SECURITY WITH THREAT INTELLIGENCE

The sophistication and complexity of attacks are increasing every day, as hackers develop new tools and discover new vulnerabilities to exploit. To keep up with attackers, security vendors must have granular visibility into emerging threats as they are developing anywhere in the world. In addition, vendors need the capability to quickly develop new rules to mitigate emerging threats and push them into global application deployments.

Because of the global scale of the Akamai Intelligent Platform, Akamai has unmatched visibility into

Internet activity and active attacks. Akamai is able the rapidly analyze the vast amounts of data aggregated by our platform and services to improve security for our customers by:

- Identify new attack trends as they develop or new attack vectors as they are first used.
- Proactively warn at-risk customers of an emerging threat or adjust the security posture of protected websites and other Internet-facing applications.
- Develop WAF rules to mitigate newly discovered attack vectors while refining existing ones to improve the accuracy of our protection against web application attacks.
- Improve the tools and processes utilized by Akamai's global SOC to detect, identify, and mitigate future attacks more quickly and effectively.
- Issue specific threat advisories to customers through Akamai's threat intelligence services.

ACCREDITATION, COMPLIANCES, & ACCEPTABLE USE

Akamai services currently meet following compliances and accreditations:

- FedRAMP- Provisional JAB ATO
- FISMA
- FIPS 199- Low and Moderate ATOs
- PCI- Tier 1 Merchant Service Provider
- ISO-27002

In addition, Akamai's solutions are complaint with several key Office of Management and Budget (OMB) mandates; enabling federal government customers to meet and maintain OMB compliances, including:

- **IPv6** as per the requirements of OMB Mandate regarding Transition to IPv6, dated September 28, 2010
- **HTTPS** as per OMB Mandate M-15-13
- **DNSSEC** as per OMB Mandate M-08-23

FedRAMP- Provisional JAB ATO

The Akamai Intelligent Platform has been FedRAMP accredited with a Joint Authorization Board (JAB) Provisional Authority to Operate (P-JAB-ATO) that meets the FedRAMP requirements as a Public Cloud Service Model and an Infrastructure as a Service (IaaS) Model. A FedRAMP P-JAB ATO is a certification by the JAB of Akamai's compliance with FISMA as well as several of the NIST Special Publications, including NIST 800-53, and FIPS publications. Akamai's FedRAMP information can be found here.

The accreditation boundary for the Akamai Content Delivery Network (CDN) covers a majority of Akamai's infrastructure and services, including:

Content Delivery Infrastructure & Services		Internal Systems & Infrastructure
HTTP Delivery Edge Servers	Luna Control Center Portal	Akamai NOCC
DNS & DNSSEC Service	HTTPS (Secure Delivery) Edge servers	KMI
Streaming Servers	Global Traffic Management (GTM) System	Akamai's DNS Servers
NetStorage		

FISMA

Prior to our FedRAMP Accreditation, Akamai met FISMA compliance and supported a Moderate SP 800-53 baseline of controls with one sub-network supporting a High controls baseline with additional

configurations and products prior to our FedRAMP accreditation in August 2013.

FIPS

Akamai had received an Authorization to Operate (ATO) from the US Department of Homeland Security (DHS) and the Nuclear Regulatory Commission at a FIPS-199 LOW as well as an ATO as part of a larger US Air Force system at a FIPS-199 Moderate. Our DHS ATO and support package is available as a reference, from our DHS COTR upon request.

PCI- Tier 1 Merchant Service

Akamai is certified as a Tier 1 Merchant Service Provider under PCI-DSS, as evidenced by our listing on the Visa website. To maintain this accreditation, Akamai undergoes quarterly network scanning and annual penetration tests by a 3rd party certified by the PCI standards council.

ISO-27002

Akamai has implemented an Information Security Management System (ISMS) based on the ISO 27001/2 (formerly 17799) Code of Practice for Information Security Management and undergoes an annual assessment by an independent third party in accordance with the ISO standard.

Code of Ethics & Regulatory Compliance

Akamai is committed to conducting our business with the highest level of ethics and integrity and in compliance with all applicable laws and regulations. Akamai expects all of our executive officers and managers to be leaders in adhering to high ethical standards and all employees to follow suit. We strive to deal honestly and fairly with all parties with whom we interact in the course of our business. To formalize this commitment, Akamai has adopted a Code of Business Ethics that applies to all of our employees. In addition, Akamai adheres to Sarbanes-Oxley Compliance regulations, quarterly employee trainings.

ACCEPTABLE USE POLICY

Akamai maintains a strict Acceptable Use Policy (AUP) to maintain the integrity of traffic delivered on its network. Akamai will not allow any customers to leverage any of our services to engage in criminal or malicious activity. This includes delivering or protection of sites which engage in or support illegal, or nefarious activities, such as:

- Lunching of cyber attacks
- Hactivistism
- Phishing and malware distribution
- Terrorist propaganda

SERVICE LEVEL AGREEMENTS

All of Akamai's services include a 100% availability SLA. While other cloud service providers also have strong Availability SLAs, even 100%, often they have exclusions for high volume events, such as DDoS attacks.

Akamai's security solutions include specific Time to Mitigation (TTM) SLAs. TTM is not the time it will take to begin mitigation, rather the time it will take to complete mitigation.

Finally, Akamai's application layer security solutions include a performance improvement SLA.

Appendix 2: Akamai's Participation in Industry Organizations

SECURITY & STANDARDS GROUPS

Akamai actively participates with a range of industry working groups. The participations is through direct membership or individual contributions to the specific Groups.

Specific material on Akamai's Security Compliance for ISO, HIPAA, SOX, and PCI DSS can be found here:

https://www.akamai.com/us/en/our-thinking/information-security/compliance/

SECURITY GROUPS

A Safe and Security network is critical to Akamai's success. Akamai has corporate and individual participation in all the critical industry security activities. This includes formal and informal (investigative) security communities.

FS-ISAC - (Financial Services - Information Sharing and Analysis Center)

https://www.fsis ac.com/ (Company Membership)

OWASP - Open Web Application Security Project

https://www.owasp.org/ (Individual Contributions through local chapters)

US CERT's NCCIC (National Cybersecurity and Communications Integration Center)

https://www.us-cert.gov/nccic

US NIAC (National Infrastructure Advisory Council)

https://www.dhs.gov/national-infrastructure-advisory-council (Working Group Participation)

US NSTAC (National Security Telecommunications Advisory Committee)

https://www.dhs.gov/national-security-telecommunications-advisory-committee (Working Group Participation)

US NCC - National Coordinating Center for Communications (NCC)

https://www.dhs.gov/national-coordinating-center-communications (Full Member)

DNS-OARC - DNS Operations, Analysis, and Research Center

https://www.dns-oarc.net/ (Membership. DNS-OARC is the active DNS Operations and "DNS CERT")

(ISC)^2 - International Information System Security Certification Consortium

https://www.isc2.org/cissp/default.aspx

Singapore (CSA) - Cyber Security Agency

https://www.csa.gov.sg/ (Corporate and Individual Participation in Working Groups)

FCC CSRIC - Communications Security, Reliability and Interoperability Council

http://transition.fcc.gov/pshs/advisory/csric/ (Akamai appointed membership to the Council. Active

participation in CSRIC Working Groups)

FIRST - Forum of Incident Response and Security Teams

https://www.first.org/ (Individual participation with FIRST Working Groups and Conferences).

STANDARD ORGANIZATIONS

Akamai's industry innovate requires deep collaboration in standards groups in many parts of the world. Contribution to the standard groups are through active membership, working group participation, or individual participation.

IETF - Internet Engineering Task Force

http://www.ietf.org (Individuals actively participating in standards activities)

W3C - World Wide Web Consortium (W3C)

https://www.w3.org/ (Full Member)

ETSI - European Telecommunications Standards Institute

http://www.etsi.org/ (Participant in ETSI Work Groups)

PCI Security Standards Council

https://www.pcisecuritystandards.org (Full Membership)]

INTERNET GOVERNANCE

Governance of the Internet keeping faith to the roots of the Internet is in Akamai's critical interest. As such, Akamai has representatives participating in key Internet Governance and policy working groups.

ICANN - Internet Corporation for Assigned Names and Numbers

https://www.icann.org/ (Individual Working Group participation and G-TLD)

NETWORK OPERATOR GROUPS

Network Operations Groups (NOGs) have been critical to keep the Internet Operational, Respond to Security Incident, and Exploring Options for Scaling. The Internet did not build itself. People build the Internet. Akamai has been in the middle of these "NOGs," being active participants and frequent speakers. NOG meeting are critical for network professionals to meet, teach each other, argue, come to agreement, and solve really tough problems that impact everyone on the Internet.

Asia Pacific

- APNIC – Asia Pacific Network Information Centre
- APOPS – Asia Pacific OperatorS Forum
- APRICOT – Asia Pacific Regional Internet Conference on Operational Technologies
- AUSNOG – Australian Network Operators Group
- BDNOG – Bangladesh Network Operators Group
- CNNOG – China Network Operations Group
- HKNOG – Hong Kong Network Operators Group
- IDNOG – Indonesia Network Operations Group

- JANOG – Japanese Network Operators Group
- MYNOG – Malaysian Network Operations Group
- NZNOG – New Zealand Network Operators Group
- PANOG – Pakistan Network Operations Group
- PACNOG – Pacific Network Operators Group
- SANOG – South Asian Network Operators Group
- SGNOG – Singapore Network Operations Group

Africa

- AfNOG – Africa Network Operators Group
- CMNOG – Cameroonian Network Operators' Group
- NGNOG – Nigerian Network Operators' Group
- SAFNOG – South African Network Operations Group
- SdNOG – Sudan Network Operator's Group

Americas

- ARIN, the American Registry for Internet Numbers
- CaribNOG – Caribbean Network Operators Group
- ISPCON, an Internet Service Providers' Convention.
- LACNIC, the Latin American and Caribbean IP address Regional Registry
- LACNOG – Latin American and Caribbean network operators group
- NANOG – North American Network Operations Group
- NOGCHILE – Network Operations Group for Chile

Europe

- BENOG – Belgian Network Operators Group
- DENOG – German Network Operations Group
- DKNOG – Danish Network Operators Group
- EOF – the European Operators Forum WG
- ENOG – Eurasia Network Operators Group
- ESNOG – Grupo de Operadores de Red Españoles (Spain)
- FrNOG – the French Network Operators Group
- GTER – Grupo de Trabalho de Engenharia e Operacao de Redes
- INOG – Irish Network Operators Group (iNOG)
- ITNOG – ITalian Network Operators Group
- NLNOG – Ring: Netherlands Network Operations Group
- NordNog – the Nordic Operator Group
- PLNOG – Polish Network Operations Group
- RIPE and the RIPE Network Coordination Centre, which promote collaboration among wide-area network operators in Europe.
- SwiNOG – the Swiss Network Operators Group

- TRNOG – Türkiye Network Operatörleri Grubu (Turkey)
- UKNOF – United Kingdom's Network Operator's Forum

Middle East

- MENOG – Middle East Network Operations Group

PEERING FORUMS

Peering Forums are a specialized outgrowth from the Peering Workshops @ NANOG, RIPE, and APRICOT. Peering forums focus conversation on the "interconnection" between Communication Service Providers (CSP)s, Internet Service Providers (ISPs), and CDN Operators. Akamai representatives attend, actively participate, present, and interact at these Peering Forums.

Global Peering Forum (GPF)

European Peering Forum (EPF)

Equinix Asia Peering Forum (APF)

Equinix Japan Peering Forum (JPF)

[Volume 4 / Number 1]

akamai's [state of the internet] / security

Q1 2017 report

AT A GLANCE

Web application attacks, Q1 2017 vs. Q1 2016
35% increase in total web application attacks
57% increase in attacks sourcing from the U.S. (current top source country)
28% increase in SQLi attacks

Web application attacks, Q1 2017 vs. Q4 2016
2% decrease in total web application attacks
20% increase in attacks sourcing from the U.S. (still top source country)
15% decrease in SQLi attacks

DDoS attacks, Q1 2017 vs. Q1 2016
30% decrease in total DDoS attacks
28% decrease in infrastructure layer (layers 3 & 4) attacks
19% decrease in reflection-based attacks
89% decrease in attacks greater than 100 Gbps: 2 vs. 19

DDoS attacks, Q1 2017 vs. Q4 2016
17% decrease in total DDoS attacks
17% decrease in infrastructure layer (layers 3 & 4) attacks
14% decrease in reflection-based attacks
83% decrease in attacks greater than 100 Gbps: 2 vs. 12
**Note: percentages are rounded to the nearest whole number.*

What you need to know

- Reflection attacks continued to comprise most DDoS attack vectors and accounted for 57% of all mitigated attacks.
- "DNS Water Torture Attacks," a DNS query flood included in Mirai malware, targeted Akamai customers in the financial services industry. Details are provided in this quarter's ***Attack Spotlight.***
- Akamai welcomes Wendy Nather, Sr. Security Strategist from Duo Security, as the first Guest Author.

LETTER FROM THE EDITOR / The *Q1 2017 State of the Internet / Security Report* represents analysis and research based on data from Akamai's global infrastructure and routed Distributed Denial of Service (DDoS) solution.

Technology milestones are often marked by a significant event, followed by a long adoption phase. When referring to consumer adoption of technology, this is called the "hype cycle," a term created by the consulting firm Gartner. The initial hype surrounding a product far exceeds its capabilities in the real world, followed by a period of disillusionment and a slow integration into the fabric of our lives. The world of DDoS attack tools differs little from other technologies; new tools used by attackers follow a similar cycle of hype and integration. However, DDoS technology acceptance often proceeds at a much faster pace than consumer technologies, as there is much less resistance to change within the relatively small community of malicious actors.

As shown over the last half year, the Mirai botnet is an example of a disruptive technology working its way through the cycle. The development of Mirai happened quietly behind the scenes, while the first round of DDoS attacks were startling in their size and capability. The botnets' capabilities quickly moved into a stage where contention for Internet of Things (IoT) devices reduced the size of attacks considerably. While many of the largest DDoS attacks observed this quarter were still based on Mirai-derived botnets, they were not as large as the initial attacks. What follows is the integration of the use of IoT as another part of the fabric of DDoS botnets and malware.

As we discussed in last quarter's report, there were long-term consequences to the release of Mirai. First, competitive forces drove botnet herders to keep up with Mirai's technology or risk losing market share. The creators of other botnets are working to generate comparably-sized attacks.

Secondly, other botnets families, such as BillGates, started adding new features, some taken directly from leaked Mirai source code. Meanwhile, Mirai has continued to splinter and evolve. There is now a variant which infects Windows systems, not to recruit them as attack nodes for the botnet, but to further expand the botnet by scanning and infecting Linux devices.

This quarter's Attack Spotlight includes our research into one of the Mirai DDoS tools used against financial services organizations. Called "DNS Water Torture" in Mirai's code, this DNS query flood generates relatively limited volumes of traffic, but can create denial of service outages by consuming the target domain's resources in looking up randomly generated domain names in great numbers. Each query ties up memory and processor cycles, preventing the target from processing legitimate traffic.

We also observed a new reflection attack vector, Connectionless Lightweight Directory Access Protocol (CLDAP). At this point, the protocol has not been a significant source of attack traffic, but the lack of contention for the resource could change its popularity. A link to the threat advisory is provided in ***Cloud Security Resources***.

We are pleased to host a guest author this quarter: Wendy Nather, Principal Security Strategist at Duo Security. See what she has to say about the challenges of managing corporate security, given the current state of the Internet.

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***. The views of Ms. Nather are her own and do not necessarily reflect the opinions or perspectives of Akamai.

The state of the Internet is...complicated, as always.

Consider these changes over the past decade:

Corporate and Consumer Use Are Intertwined / It used to be that you went to work in the office, used corporate software, and then went home and used completely different software on your home computer. Now, more often than not, you've got a corporate login and a personal login with the same SaaS provider and you're using the same apps on your phone (Gmail, Dropbox, LastPass, etc.). Unless you're working in a strictly segmented environment, the expectation is that you'll be using applications for both purposes and alternating at the drop of a hat, regardless of which network you're currently connecting to.

BYODon't / Some organizations have embraced the use of personal devices, and others haven't, but it's becoming harder to enforce a "no BYOD" policy when both the endpoint and the resources they're accessing are outside of the corporate perimeter. Unmanaged personal devices raise the specter of risks ranging from unpatched vulnerabilities to e-discovery requirements that include searching your employees' phones. And that's not even counting wearables and other Things.

Password Policies / Remember when you only had a dozen usernames and passwords? Yeah, neither do I, and here we are. A typical online user could have literally hundreds of online accounts, some of which predate today's password managers. Under pressure from bulk credential theft and compliance requirements, every system owner is being driven to require longer, more complicated and unique passwords. But the days of password rules such as "upper and lower case, one number, one special character, two emojis, and a squirrel noise" are going to come to an end; users are going to push back as soon as the absurdity becomes clear. Ubiquitous, consistent, and usable password managers are going to have to evolve into an application interface to shield everyday people from the malignant growth of complex passwords.

Guest Author

Wendy Nather
Principal Securty Strategist
DUO Security

To Sum Up / Our interaction with the Internet has evolved to "anytime, anywhere, using any device and software, for any purpose." That means that enterprises have to address the security issues in ways that don't rely exclusively on traditional boundaries ("our network," "our software," "our hardware"). And they have to be able to distinguish business data from personal data, which were created at the same time of day, in the same location, using the same applications, and stored in the same formats on the same hardware and services. Users expect a seamless experience that doesn't require them to sacrifice a chicken every time they switch between corporate and personal contexts — and they deserve one.

The identity is the new boundary, together with the context. When you log into Gmail with your personal credentials, you're in charge of the security requirements you set for accessing your data; when you use your corporate credentials to log in, your employer has to specify what's required to access business data, such as the combination of username, password, other authentication factors, and managed device. It's the same service, the same software, and the same person, but there are different stakeholders based on the ownership of the data.

Adapting to this new boundary, Google built a framework for their internal use and dubbed it ***BeyondCorp***; whether they're calling it "***zero-trust***," or "perimeterless," many organizations are looking to adopt it ***in their own ways***. The important point is that the security shouldn't rely solely on the traditional perimeter, and should accommodate the needs of both the user and the enterprise.

Putting the user on equal footing with the data owner is a welcome trend, and it's one that holds great promise for the ongoing challenge of securing the Internet.

[SECTION][1] EMERGING TRENDS

The median size of DDoS attacks has fallen steadily since the beginning of 2015. At the beginning of 2015, the median DDoS attack size was 4 Gbps. Two years later, at the beginning of 2017, the median attack size was just over 500 Mbps. Not to say huge attacks aren't happening — mega attacks topping 100 Gbps occur every quarter — but half of all attacks are between 250 Mbps and 1.25 Gbps in size. Even these smaller attacks can harm unprepared organizations. Web application attacks shifted subtly towards the U.S. this quarter, both as a source and as a target. This type of attack is important not because of their size, but because they attack the underlying fabric of sites, either tying up resources or pulling information from the database powering sites.

The impact of IoT devices and dozens of attacks from the Mirai botnets since last September has had a strong practical effect on the security needs of organizations. The mega attacks are outliers that represent the limits enterprises must be prepared to defend against. However, the overwhelming number of smaller attacks means that these mega

attacks have little impact on the trend lines that define the median attack size, which is a better indicator of what an organization is most likely to see.

The majority of attacks are still small relative to the largest Mirai attacks, but they don't need to be big to be effective. If we consider that many businesses lease uplinks to the Internet in the range of 1–10 Gbps, any attack exceeding 10 Gbps could be "big enough" and more than capable of taking the average unprotected business offline.

At the same time, the effects of IoT are not to be underestimated, and the IoT ecosystem has drawn the attention of a wider audience. A recent example is malware that compromises Internet-enabled toasters to mine Bitcoins[1], an effort that appears to have been an ineffective proof of concept. Another trend is represented by the BrickerBot botnet, which attacks systems exposed directly to the Internet with default Telnet passwords apparently in an attempt to prevent their use by the Mirai botnet. If this botnet is unable to disconnect the target device from the Internet, it corrupts the configuration, permanently bricking the devices[2]. Neither of these examples are major threats, but they do show a significant increase in attention from both the hacker and security communities.

There is one factor that seems to be affecting the DDoS landscape as a whole: law enforcement. Early attacks by the Mirai botnets appear to have been triggered by the announcement of the arrests of two teens in Israel who were responsible for the vDos botnet[3] — a DDoS-for-hire tool that netted them hundreds of thousands of dollars. More recently, Europol coordinated the arrest of 34 individuals across 13 countries as part of an effort called Operation Tarpit[4]. Operations like Tarpit target the largest services responsible for DDoS attacks directed at banks, gaming companies, and retailers. This can have a significant effect in reducing the number of attacks on these organizations.

Despite the overall reduction in volumetric DDoS attacks, Akamai has seen a significant increase in the amount of traffic in reflection attacks. Taking advantage of the nature of DNS, NTP, and other protocols, attackers make seemingly legitimate requests of servers, causing them to spew traffic at the attacker's true target. Akamai recently released a threat advisory about adding a new DDoS reflection source, CLDAP[5]. Reflection attacks are much more difficult to track back to the botnets that originate the attacks.

In all likelihood, DDoS attacks will increase in size and frequency. We anticipate more frequent small-scale attacks, but the largest attacks will almost certainly continue to grow. As previously noted, we expect mega attacks to continue to have an outsized impact on DDoS trends in the coming years.

[SECTION]^2
DDoS ACTIVITY

2.1 / DDoS Attack Vectors / As the research team dove into early 2017 data, we first examined infrastructure-related attack data. Invariably, infrastructure attacks are the largest component of our quarterly volumetric attack data. In Q1, these attacks accounted for roughly 99% of the overall attack traffic. That's likely because it's trivial for an attacker to launch a volumetric attack in comparison to the technical understanding needed to make effective use of application layer tools.

Application layer DDoS attacks such as GET, PUSH, and POST floods remained a small component of the overall DDoS attack landscape. Two years ago, in Q1 2015, application layer DDoS attacks accounted for 9% of all attacks. In Q1 this year, only 0.6% of DDoS attacks targeted the application layer. Most application layer attacks aren't designed for denial of service.

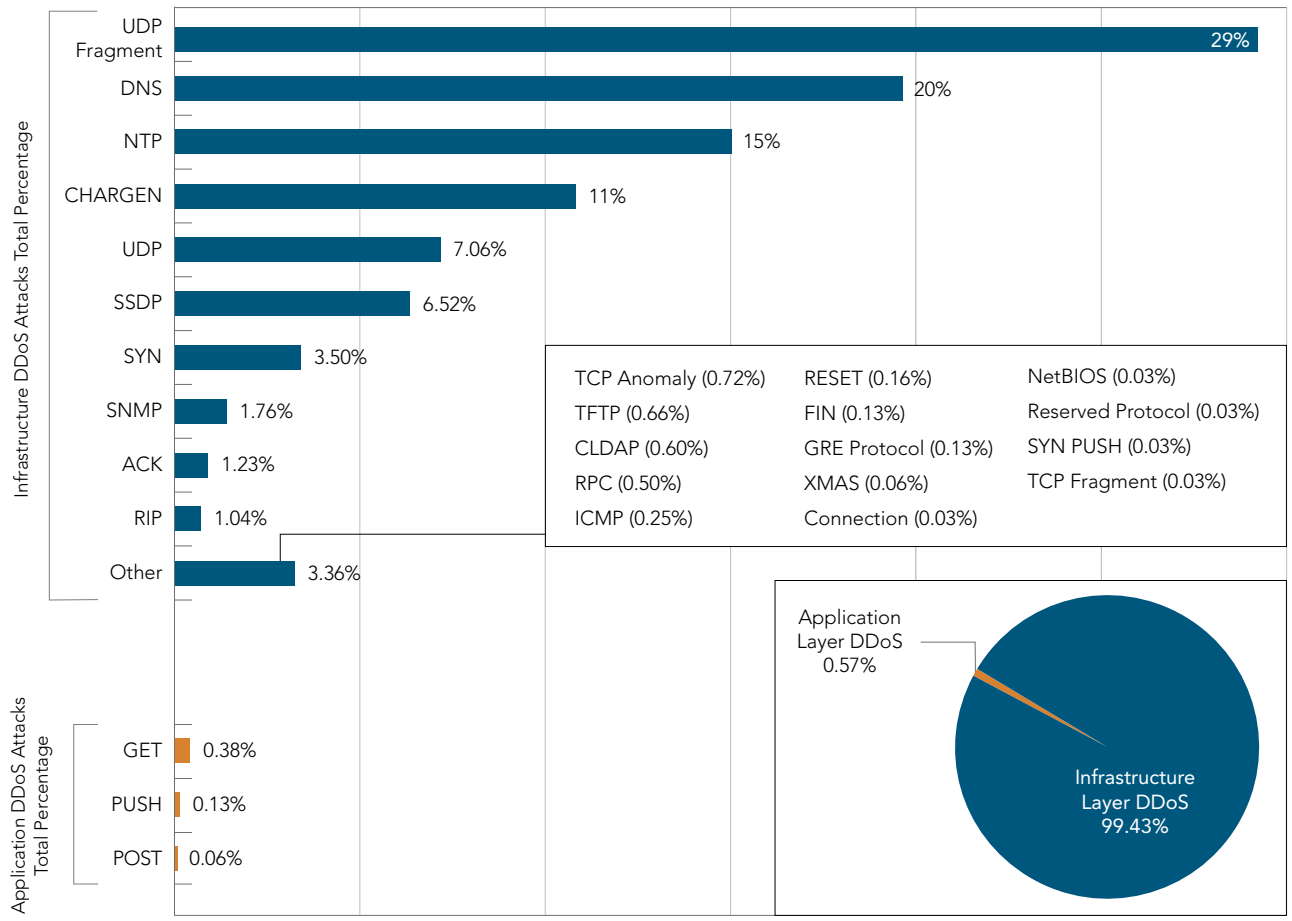

Figure 2-1: UDP fragment, DNS, and NTP continued as the top three DDoS attack vectors, while reserved protocol floods and connection floods made rare appearances in the attack vectors list

The top four infrastructure DDoS related attacks were the same as in recent quarters. UDP fragments, DNS floods, NTP floods, and CHARGEN attacks dominated, as shown in Figure 2-1. UDP fragment, NTP, and CHARGEN rose compared to the previous quarter, while DNS attack traffic fell slightly from 21% to 20%.

Organizations can keep their servers from participating in these DDoS attacks if they ensure that services such as CHARGEN and NTP are either not accessible from the Internet or are patched. Older NTP daemons, as an example, send large amounts of reflected traffic at the intended attack target in response to relatively small illegitimate requests from attackers. This traffic amplification factor is one reason why attackers continue to use NTP reflection even as fewer and fewer unpatched NTP servers remain on the Internet. One easy fix is to confirm the NTP daemons that are running in your environment are well patched. No defender wants to make the job of an attacker easier.

DDoS attacks are an ever present danger and it's important that defenders make sure that they are practicing proper security hygiene to avoid inadvertently participating in attacks. It is essential to ensure that services such as CHARGEN and NTP are patched and firewalled off where they are not required to be available to the wider Internet.

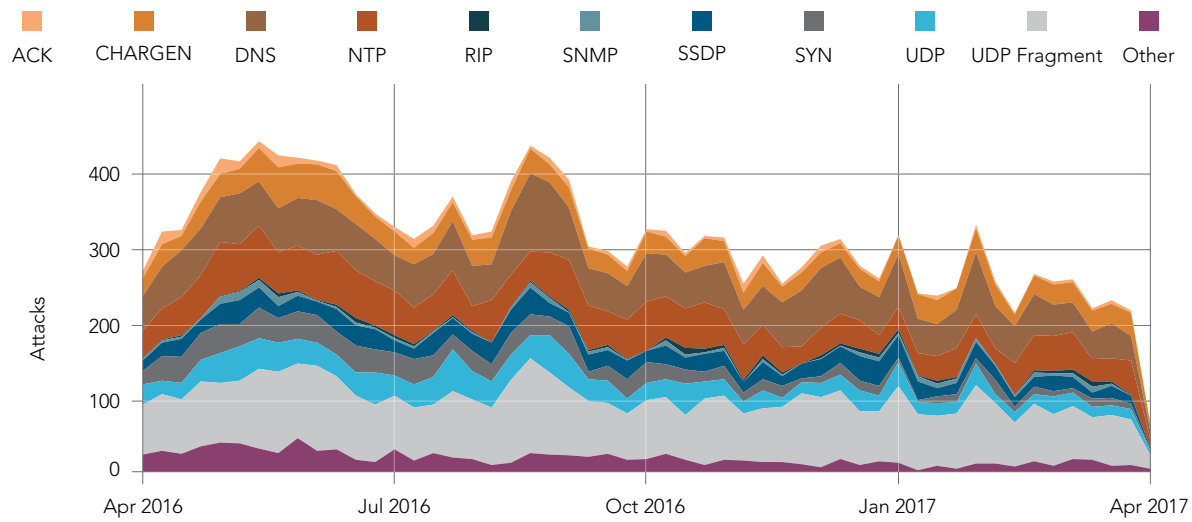

Figure 2-2: Attack traffic for the 10 most frequent attack vectors shows reflection attacks, such as NTP and CHARGEN, continue to generate large amounts of DDoS traffic

In looking at the 10 most frequent attack vectors per week, we see ACK, CHARGEN, and DNS in the top three, with NTP taking fourth place in the list.

One item of note, that's unfortunately consistent, is the presence of CHARGEN on the list. CHARGEN traffic rose to 11% of DDoS attack traffic in Q4, up from 8% in the previous quarter. This protocol is used as a diagnostic port on printers and this service should not be exposed to the Internet at large.

The percentage of the Internet attack traffic related to NTP was relatively flat this quarter; the .5% change in traffic is well within our margin of error. This attack vector can be utilized by attackers to amplify their DDoS attacks. It is not outside of the realm of possibility to posit that this will result in a correlation with the rise of IoT-related botnet platforms — the rationale being that it will only be a matter of time before attackers can implement this in their platforms.

Several individuals from some of the criminal organizations responsible for the day-to-day operations and upkeep of these attack platforms have been incarcerated. Incarcerations alone may not limit the number of attacks in the long term as other operators will likely fill the void. This is especially true when one considers that there is money to be made from facilitating these attacks as a service offering.

2.2 / Mega Attacks / The mega attacks — those over 100 Gbps — were in shorter supply in the first quarter of 2017. While this may result in a drop in the number of attacks, the reduction could be short-lived. Several large DDoS crews were arrested in the waning days of 2016, which could be linked to the drop in mega attacks.

Another contributing factor to the drop in large attacks could be the evolving use cases for botnets like Mirai. As an example, attackers have created a proof of concept that uses the Mirai botnet for Bitcoin mining[6]. While this activity might seem clever on the surface, there's little benefit to the attackers; the IoT devices employed by the Mirai botnets do not have the requisite computing power to mine Bitcoins effectively. Despite the botnet being an inefficient Bitcoin mining tool, this may be an indicator that Mirai and other botnets may be used for a diverse set of purposes in the future.

2.3 / Attack Spotlight: Mirai DNS Water Torture Attack Summary / Akamai observed a series of DDoS attacks leveraging the Mirai DNS Water Torture Attack. DDoS attacks using this DNS query vector were first observed on Jan. 11, 2017, targeting several Akamai customers in the financial services industry. The attack activity began with five consecutive days of attacks, followed by a four-day reprieve before concluding with a final attack on Jan. 20. Aside from UDP and TCP attacks observed on Jan. 12, all the other attacks were Mirai DNS query floods.

Payload Samples / The Mirai DNS Water Torture Attack follows normal DNS recursion paths. As a result, the attacker cannot select a specific IP address at the target site.

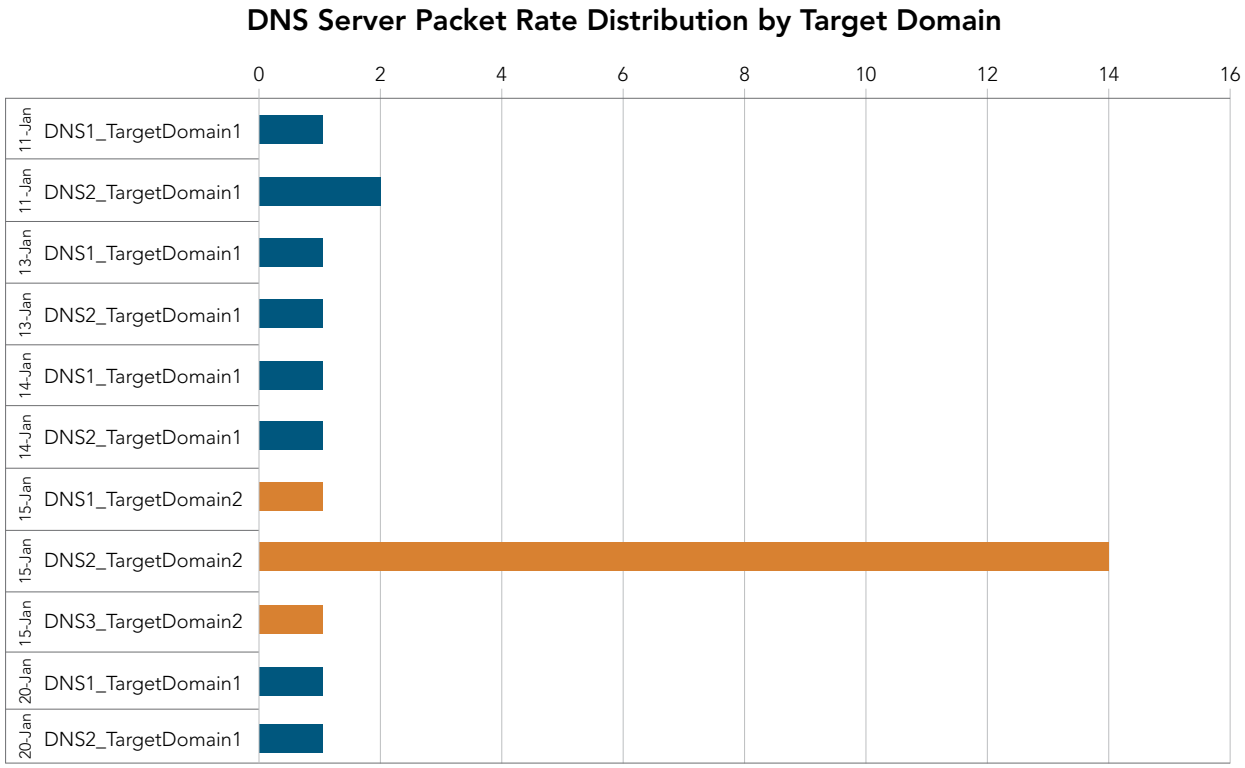

Figure 2-3: Peak packet rates observed on DNS servers receiving Mirai DNS attacks reached a high of 14 Mpps on Jan. 15, 2017

Most of the DNS servers received queries at a fairly even rate during the attack, with the exception of an attack observed on Jan. 15, when one of three DNS servers received 14 Mpps of attack traffic, as opposed to the 1-2 Mpps other DNS servers received. The queries observed during these attacks aligned with the Mirai DNS Water Torture Attack.

The sample payload signatures in Figure 2-4 represent a flood of queries, each containing a random 12-character subdomain string. The IP addresses and targeted domains have been redacted.

```
DNS Query Flood (Mirai DNS Water Torture Attack)
08:10:13.574610 IP x.x.x.x.47565 > x.x.x.x.53: 10077 [1au] A? e4hob2e7w1t7.<redacted>. (xx)
08:10:13.591581 IP x.x.x.x.52465 > x.x.x.x.53: 15764 [1au] A? sjjbm0s2ov00.<redacted>. (xx)

06:50:44.189382 IP x.x.x.x.49326 > x.x.x.x.53: 63481% [1au] A? io1f786uo3bd.<redacted>. (xx)
06:50:44.189429 IP x.x.x.x.40566 > x.x.x.x.53: 12345% [1au] A? 0hagnikgj2vq.<redacted>. (xx)

11:14:10.707489 IP x.x.x.x.37569 > x.x.x.x.53: 25550% [1au] A? 1hartrmnaiew.<redacted>. (xx)
11:14:10.709341 IP x.x.x.x.22945 > x.x.x.x.53: 31835% [1au] A? c7wnmqek2eww.<redacted>. (xx)

04:56:19.326305 IP x.x.x.x.4210 >  x.x.x.x.53: 47369% [1au] A? lmjtjgfh7b6j.<redacted>. (xx)
04:56:19.326315 IP x.x.x.x.36408 > x.x.x.x.53: 36684% [1au] A? 2vfedrv6aha5.<redacted>. (xx)

11:48:43.171738 IP x.x.x.x.47645 > x.x.x.x.53: 59218 [1au] A? 02uqhuovfi1f.<redacted>. (xx)
11:48:43.171749 IP x.x.x.x.47371 > x.x.x.x.53: 62949 [1au] A? qo5etoh5foab.<redacted>. (xx)
```

Figure 2-4: Payload of DNS query flood, called the Mirai DNS Water Torture attack, with the target domain names redacted

On Jan. 12, malicious actors changed tactics. After a day of DNS query floods, the attackers began attacking a DNS server directly with a UDP flood, as shown in Figure 2-5. They also made use of one of Mirai's TCP flood attacks on TCP port 443, a port commonly used for transmission of encrypted web traffic. This type of Mirai attack is called Mirai TCP STOMP.

```
UDP Flood – Port 53
06:17:36.688058 IP (tos 0x0, ttl 51, id 54282, offset 0, flags [DF], proto UDP (17), length 540)
    x.x.x.x.59242 > x.x.x.x.53: 56019 stat+ [b2&3=0x1786] [2646a] [49544q] [26389n] [1379au]
[|domain]
06:17:36.688063 IP (tos 0x0, ttl 52, id 24494, offset 0, flags [DF], proto UDP (17), length 540)
    x.x.x.x.44026 > x.x.x.x.53: 55693 updateA+ [b2&3=0x4b01] [24342a] [13221q] [35165n]
[62407au] Type60358 (Class 50264)? M-^_M-sM-?M-xM-hM-^KM-bM-'M-?^V^I^YM-4TTFM-~xwy^T^IM-J^X-
a^vM-6M-g[M-^GM-UM-3a7M-^\M-CIM-5M-^L"M-^Z0~^UM-<snip>[|domain]

Push Flood (Mirai TCP STOMP) – Port 443
08:18:32.564571 IP (tos 0x0, ttl 54, id 34074, offset 0, flags [DF], proto TCP (6), length 808)
    x.x.x.x.38403 > x.x.x.x.443: Flags [P.], cksum 0x4768 (correct), seq 535625728:535626484, ack
1, win 22263, options [[bad opt]
08:18:32.564735 IP (tos 0x0, ttl 54, id 24701, offset 0, flags [DF], proto TCP (6), length 808)
x.x.x.x.38403 > x.x.x.x.443: Flags [P.], cksum 0x0dc9 (correct), seq 535887872:535888628, ack 1,
win 22263, options [[bad opt]
```

Figure 2-5: The signatures of UDP and TCP vectors used when attackers changed tactics on Jan. 12, 2017

The UDP flood was observed against two destination IP addresses, one of which was a DNS server previously under attack from the DNS query flood. The signatures contained the standard Mirai UDP flood, using 512 byte payloads; however, they first appeared to be DNS because Port 53 was used as the the target. The other signature was a PUSH Flood set to target port 443. This type of attack completes the TCP three-way handshake prior to sending a flood of padded TCP packets. The extra data padding results in higher peak bandwidth consumption with lower packet rates — in this case the attack peaked at 120 Gbps.

Conclusion / Given the risk posed by the Mirai DNS query flood attack, all DNS servers responding for a targeted domain should be protected. Some organizations may be capable of serving this malicious traffic in addition to their normal load of legitimate queries. But even in those cases, the flood of requests puts unnecessary load on DNS systems, which often run at the edge of their capabilities. In some cases an external DNS provider is required in order to have sufficient response capabilities. Even in the case of an external provider, it can make sense to have redundant providers, a point several of last year's attacks drove home.

DDoS protection should take DNS load distribution into account. Be aware that bots may cluster within regions where vulnerable devices reside. If regional balancing is in effect, the malicious traffic may not be desirably distributed during an attack. Vectors, techniques, or targets may vary throughout the DDoS campaign. Any organization could find itself under threat of DDoS, regardless of industry. Attention needs to be given to assets that could be attacked and may be vulnerable, in addition to assets that have been attacked in the past. It's best to ensure that DDoS mitigation is in place before an attack.

DNS WATER TORTURE

DNS Water Torture / Mirai has been known to produce a specific DNS query flood. Although DNS query attacks are not as common as DNS reflection attacks, this DNS query flood can potentially cause more damage than current DNS reflection attacks. If a targeted DNS server is unprepared for a sustained flood of queries with high packet rates, DNS Water Torture can lead to a denial of service for legitimate users.

How it works / The Mirai DNS query flood does not use reflection or spoofing techniques, nor does it allow attackers to specify a target IP address. Instead, it accepts a domain name as the target for a DNS cache-busting flood. A randomized 12-character alphanumeric subdomain is prepended to the target domain. The attacking bots send their queries to their locally-configured DNS servers, which are typically DNS servers at local ISPs (Internet Service Provider). The randomized sub-domain is present to ensure that no intermediate recursive DNS server would have the response for that name cached locally. Since the response cannot have been cached, every query follows the usual path until it reaches an authoritative DNS server, the real target of the attack.

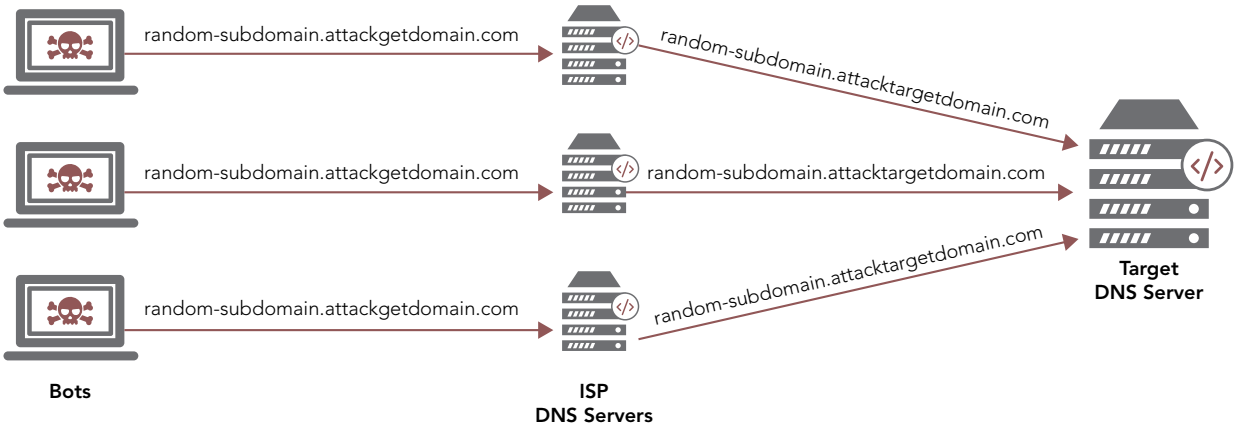

Figure 2-6: Mirai DNS attack queries are sent from bots to their local DNS servers and on to the target authoritative DNS servers

Aside from the randomized subdomain string, the queries appeared to the target authoritative DNS servers as queries from local ISP DNS servers. The full source IP addresses of the bots sending these queries were not visible.

Akamai SIRT has reproduced and tested Mirai's DNS query attack, using live malware samples from the initial documented attacks. The attack supports several customizable values as shown in Figure 2-7.

Customizable Field	Default Value	Custom Value
ToS	0	1
ID	random	1
TTL	64	123
DF	false	5
SPort	random	31337
DPort	53	8008
Domain	(user supplied)	attacktargetdomain.com
DNS ID	random	1

Figure 2-7: Customizable fields for the Mirai DNS query attack, known as the DNS Water Torture attack

Attack signatures are summarized in Figure 2-8, first with default values and then with custom values.

This attack vector was observed by Akamai SIRT in January 2017 against Akamai customers within the financial services industry.

```
Examples of DNS Parameters and Resulting Traffic:

Default DNS attack traffic with no parameters besides target domain.

00:40:40.611489 IP (tos 0x0, ttl 64, id 52446, offset 0, flags [none], proto UDP (17), length 73)
    x.x.x.x.17517 > x.x.x.x.53: 3644+ A? m3hk3nr6njv0.attacktargetdomain.com. (45)
00:40:40.611490 IP (tos 0x0, ttl 64, id 60934, offset 0, flags [none], proto UDP (17), length 73)
    x.x.x.x.43103 > x.x.x.x.53: 19269+ A? htuhwake2bkg.attacktargetdomain.com. (45)

DNS attack with all values customized.
* DNS ID value @ 0x0010 column 7, traffic shown in hex format to allow highlighting

00:48:58.620735 IP (tos 0x1,ECT(1), ttl 123, id 1, offset 0, flags [DF], proto UDP (17), length 73)
    x.x.x.x.31337 > x.x.x.x.8008: UDP, length 45
        0x0000:  4501 0049 0001 4000 7b11 7af1 c0a8 01e6  E..I..@.{.z.....
      * 0x0010:  c0a8 017a 7a69 1f48 0035 fcc2 0001 0100  ...zzi.H.5......
        0x0020:  0001 0000 0000 0000 0c6a 6976 3868 7475  .........jiv8htu
        0x0030:  6877 616b 650a 7468 652d 7669 6374 696d  hwake.attacktargetdomain
        0x0040:  0363 6f6d 0000 0100 01                   .com.....

00:48:58.620738 IP (tos 0x1,ECT(1), ttl 123, id 1, offset 0, flags [DF], proto UDP (17), length 73)
    x.x.x.x.31337 > x.x.x.x.8008: UDP, length 45
        0x0000:  4501 0049 0001 4000 7b11 7af1 c0a8 01e6  E..I..@.{.z.....
      * 0x0010:  c0a8 017a 7a69 1f48 0035 ef4c 0001 0100  ...zzi.H.5.L....
        0x0020:  0001 0000 0000 0000 0c32 626b 6733 736e  .........2bkg3sn
        0x0030:  7276 3061 730a 7468 652d 7669 6374 696d  rv0as.attacktargetdomain
        0x0040:  0363 6f6d 0000 0100 01                   .com.....
```

Figure 2-8: Attack signatures of the Mirai DNS Water Torture attack using default and custom values respectively

2.4 / Reflection Attacks / Reflection attacks continued to dominate DDoS activity. As in the previous quarter, DNS, NTP, and CHARGEN remained as the top three attack vectors, as shown in Figure 2-9. Their continued use is a symptom of subpar system and network hygiene. The steps needed to close these vulnerabilities are known and often inexpensive. The long-term health of the Internet would benefit from learning what factors lead organizations that own these systems to allow the reflection vulnerabilities to persist.

Organizations should review the scalability of their DNS infrastructure. If your primary DNS is self-hosted and it goes down, then your customers would be unable to find your website or contact you via email. Having a secondary or even tertiary DNS provider can help keep your systems available.

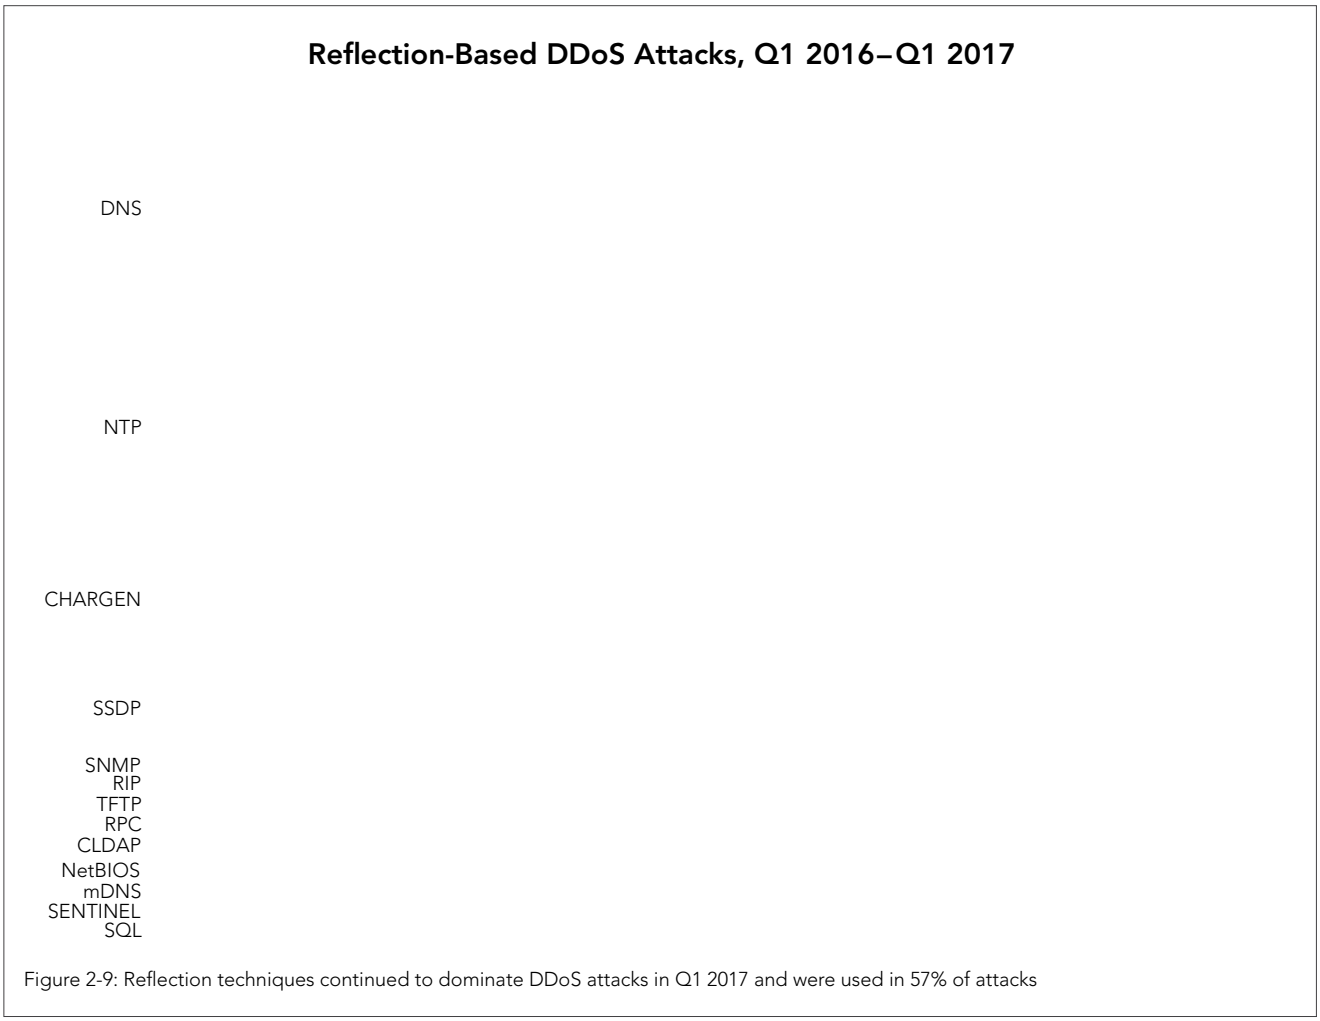

Figure 2-9: Reflection techniques continued to dominate DDoS attacks in Q1 2017 and were used in 57% of attacks

Autonomous System Numbers (ASN) designate the ISP responsible for originating the traffic and give more detail than country level statistics. Chinese ASN 4837 produced more reflection DDoS sources in Q1 2017 than the next closest ASN in Argentina. All together, the top ten reflection source ASNs accounted for 30% of the reflection DDoS sources. This analysis does not examine the density of attacks compared to the population, it shows the raw number of attacks from each ASN regardless of size.

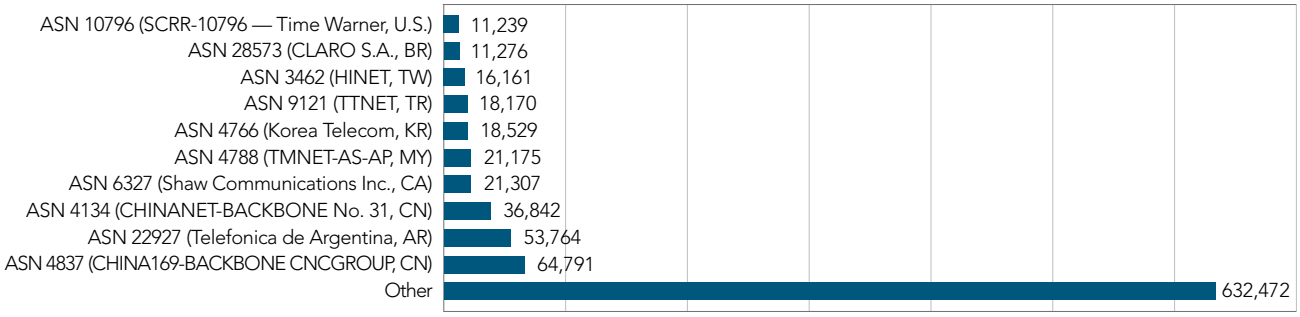

Figure 2-10: ASN 4837, in China, had the most reflection sources and ASN 22927, in Argentina, was second

The reflector data is based on observed attack sources, not the results of scans. Increased use of an attack vector can increase the number of IP addresses, especially for an attack such as Simple Services Discovery Protocol (SSDP), which uses many small devices. Use of the SSDP attack vector increased this quarter, perhaps due to attackers turning to the DDoS resources presented by IoT devices.

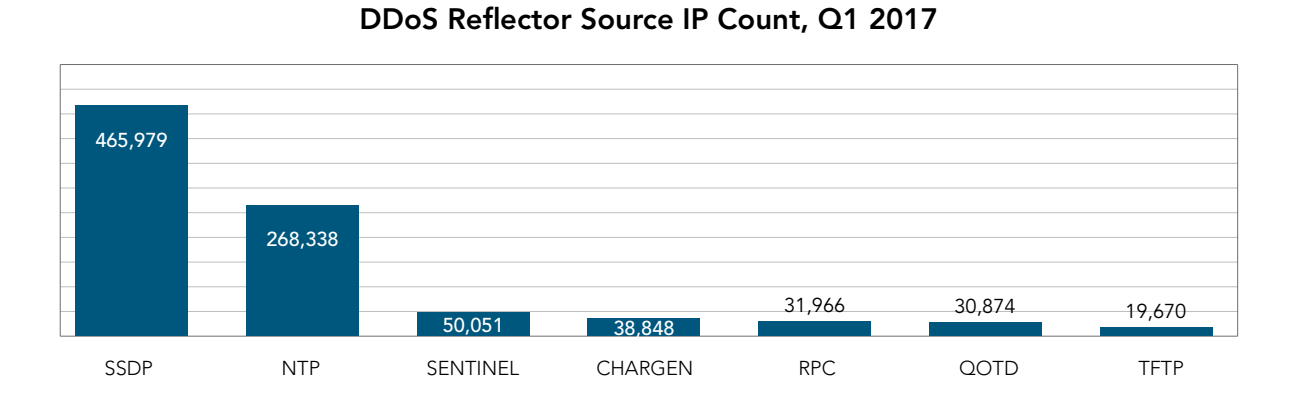

Figure 2-11: The most-used reflectors in Q1 2017 were SSDP and NTP

In Q1 2017, the use of reflectors in DDoS attacks maintained nearly the same proportions as in Q4 2016, with the notable addition of Sentinel to the top three. The number of Sentinel reflectors increased by 39% in comparison to Q4 of 2016. Sentinel reflection sources include powerful servers with high bandwidth availability, such as university servers.

SSDP reflectors continued to be the major source of DDoS reflection attacks in this quarter. The use of SSDP reflection can be directly linked to the rise of IoT botnets and the growing number of Internet-accessible consumer grade devices. These botnets are using SSDP reflectors to amplify the traffic they generate, further increasing the threat they pose.

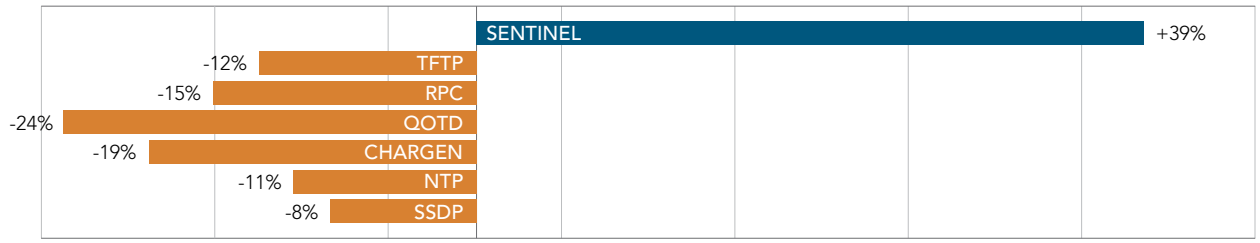

Figure 2-12: The use of Sentinel reflectors increased 39% over the previous quarter, while the use of all other reflector sources dropped 8% to 24%

Figure 2-13 shows changes in DDoS attack traffic since January 2015. The color of the dots represents how many attacks of a certain size occurred each week; the brighter colors represent a higher concentration of attacks. This plot uses a logarithmic scale, so the difference in bandwidth increases ten-fold between each major horizontal line. As a result, the attacks on the lower end of the scale appear to be more spread out, but they are actually more closely clustered numerically than attacks on the high end of the scale. The rising number of low-bandwidth attacks seen weekly is the primary reason the median size of attacks has trended downward since the beginning of 2015.

The black line represents the median (half are smaller, half are larger) attack size for each time period. In January 2015, the median attack size was 3.9 Gbps, but by the end of March 2017, the median attack size had fallen to 520 Mbps. This decline was caused in part by an overall increase in the number of weekly attacks seen by Akamai, the majority of which were smaller attacks. Growth in the number of small attacks has a more significant effect on the median than the slower growth in the number of large attacks.

The solid blue and solid red lines represent the 25th and 75thpercentile of attacks. As of March, half of all volumetric attacks seen by Akamai were between 243 Mbps and 1.3 Gbps. The dotted lines show the 5th and 95th percentiles and indicate that 90% of all attacks were between 28 Mbps and 4.8 Gbps. These ranges have long been trending closer to the median line over time, driven by an increased number of attacks since the beginning of 2015.

How does this affect enterprises? If an organization has defenses that can withstand 1.3 Gbps of volumetric DDoS attack traffic directed at its infrastructure, then it should be able to withstand 75% of current DDoS attacks. However, if the organization's uptime goals are such that it needs to withstand 95% of attacks, those defenses would need to be able to absorb an attack of 5 Gbps or more.

Even with that level of defense in place, it is important to understand that there are still a significant number of outliers — DDoS attacks generating more than 100 Gbps of traffic are common enough to be a concern.

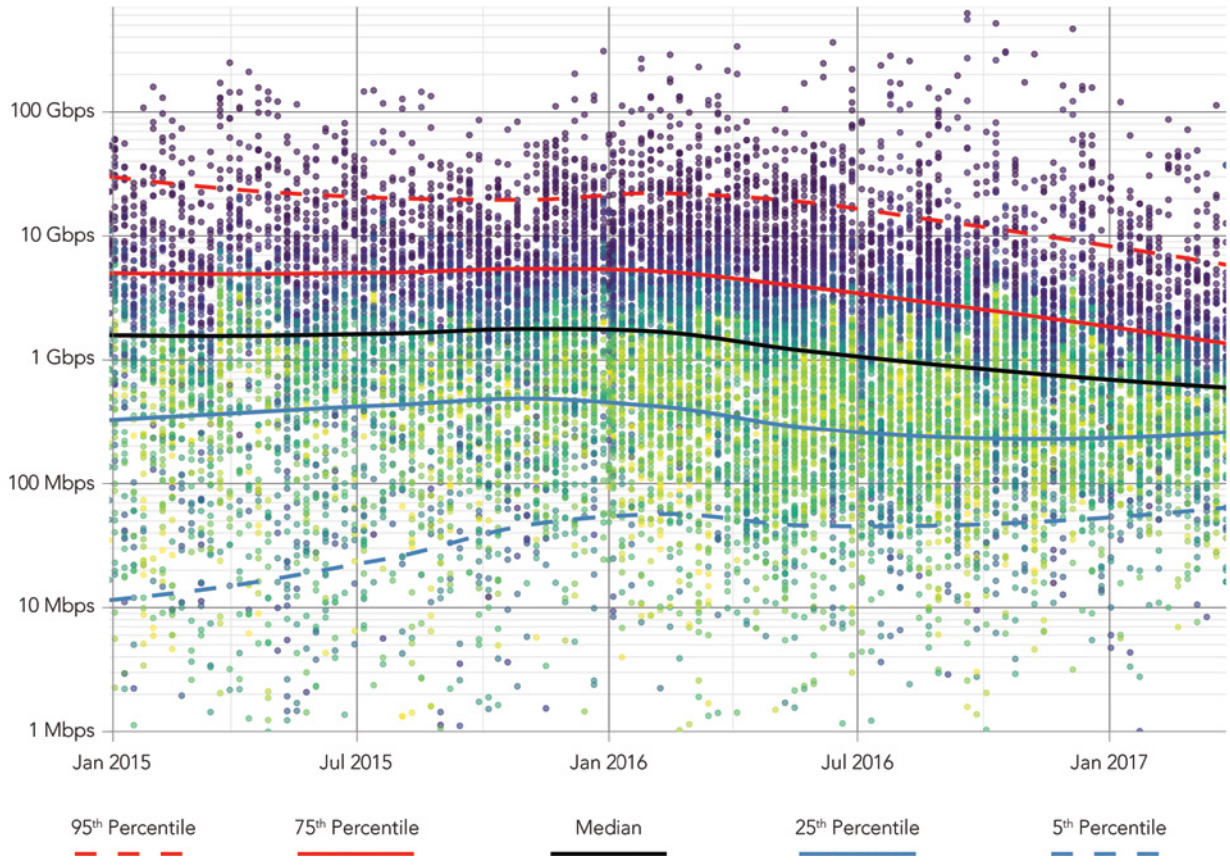

Figure 2-13: The median size of attacks has trended downward since 2016 due to an increase in the number of lower-bandwidth attacks

[SECTION]3 WEB APPLICATION ATTACK ACTIVITY

Web application vectors tend to be troublesome attack types seen across the platform. They can have a longer lasting impact than merely causing network availability outages, which we see from infrastructure-related DDoS attacks.

3.1 / Web Application Attack Vectors / We see similar patterns in the top attack types used against web applications from quarter to quarter. The top three attack vectors in Q1 of 2017 were SQLi, LFI, and XSS, as shown in Figure 3-1. These attacks continue to dominate, as they work more often than not against unprotected websites. Conversely, if your website protections are not actively blocking this sort of traffic, there is a greater risk that these sorts of attacks potentially impact your organization.

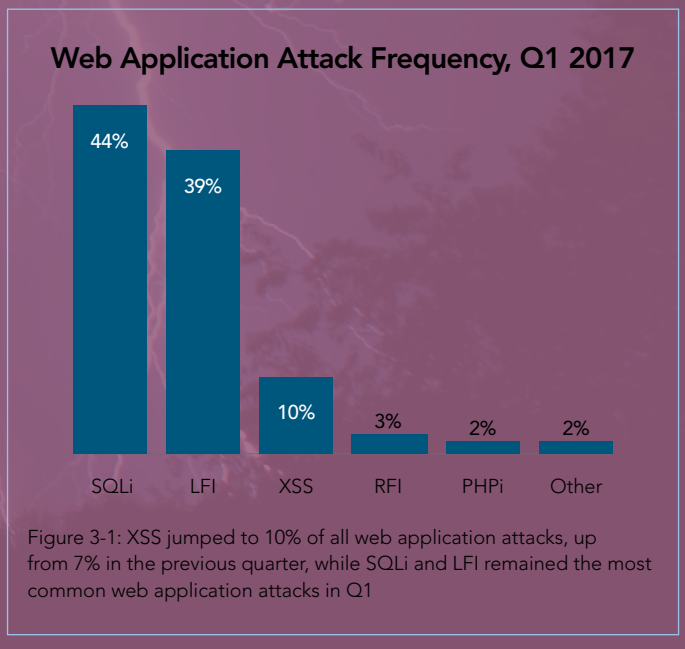

Figure 3-1: XSS jumped to 10% of all web application attacks, up from 7% in the previous quarter, while SQLi and LFI remained the most common web application attacks in Q1

3.2 / Top 10 Source Countries / The top five source countries for web application attacks in Q1 2017 were the U.S., Netherlands, Brazil, China, and Germany, as shown in Figure 3-2. For the second quarter row, Canada came in 11th place. With a small population density, it would be interesting to dig deeper into the Canadian traffic. U.S. holding on to the top position was unsurprising, but the consistent amount of attack traffic that ostensibly originated from the Netherlands is curious. This represents a large proportion of attacks from a country of only 17 million citizens. In comparison, the U.S. has just over three times the number of attacks with nearly twenty times the population.

Global Web Application Attack Source Countries, Q1 2017

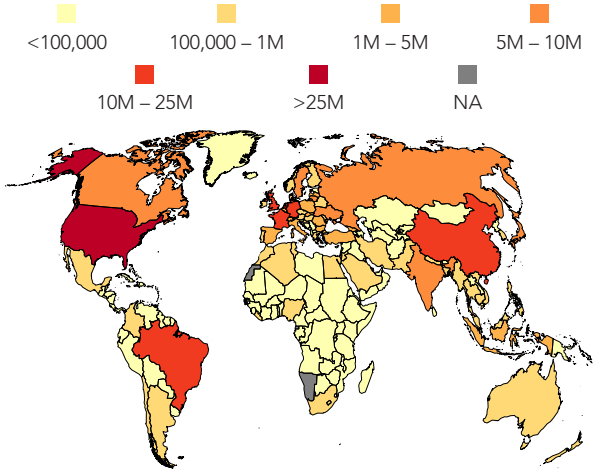

Country	Attacks Sourced	Percentage
U.S.	117,978,342	34.0%
Netherlands	43,925,118	12.7%
Brazil	28,027,292	8.1%
China	18,963,654	5.5%
Germany	16,035,037	4.6%
France	15,236,624	4.4%
U.K.	10,253,814	3.0%
Russia	8,262,292	2.4%
Ukraine	7,123,491	2.1%
Lithuania	6,566,703	1.9%

Figure 3-2: Worldwide, the U.S. continued as the top source of web application attacks at 34%, up from 28% in the previous quarter, while the Netherlands, in second, dropped to 13% from 17%

Web Application Attack Source Countries—EMEA, Q1 2017

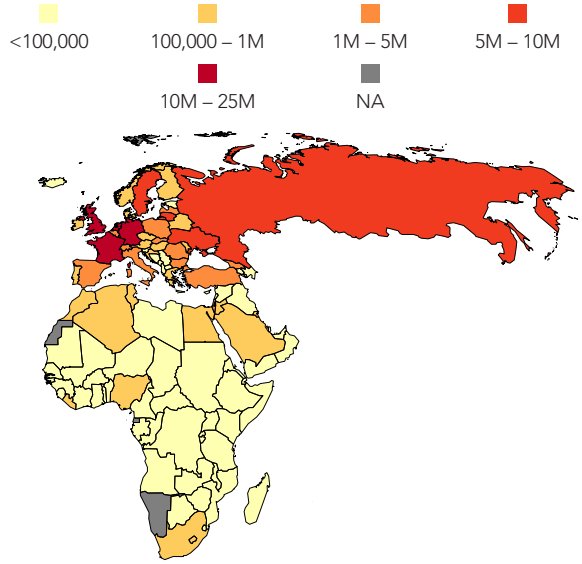

Country	Attacks Sourced	Global Rank
Netherlands	43,925,118	2
Germany	16,035,037	5
France	15,236,624	6
U.K.	10,253,814	7
Russia	8,262,292	8

Figure 3-3: The Netherlands continued as the top source for web application attacks sourced in EMEA

As demonstrated in Figure 3-5, China was the overall top source country for web application attacks in the Asia-Pacific region. Attack traffic from China increased by a third from last quarter, which cemented its place within Asia, and moved it up to fourth place worldwide.

Web Application Attack Source Countries—Americas, Q1 2017

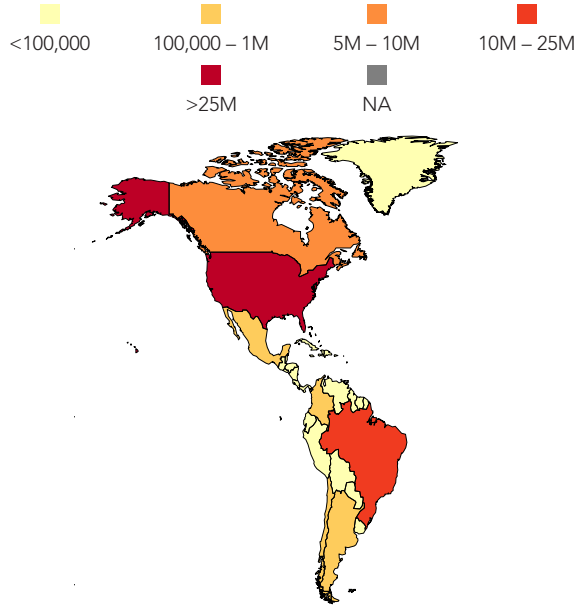

Country	Attacks Sourced	Global Rank
U.S.	117,978,342	1
Brazil	28,027,292	3
Canada	6,280,383	11
Mexico	664,477	39
Argentina	398,208	50

Figure 3-4: The U.S., Brazil, and Canada, respectively, continued as the top three sources of web application attacks in the Americas

Web Application Attack Source Countries—Asia Pacific, Q1 2017

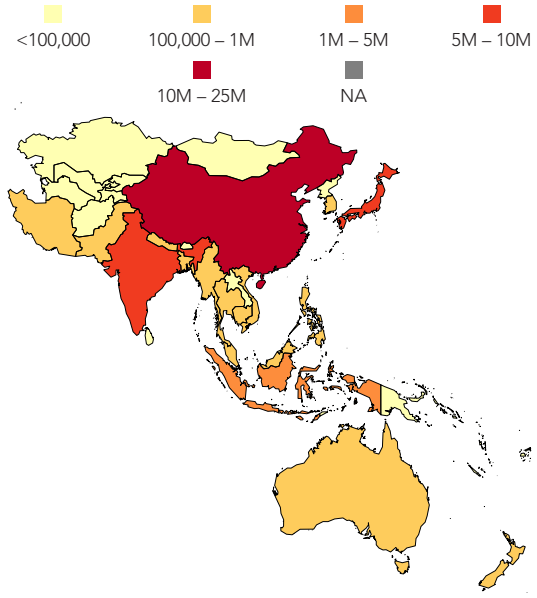

Country	Attacks Sourced	Global Rank
China	18,963,654	4
India	6,150,881	12
Japan	5,839,869	13
Singapore	4,285,527	15
Indonesia	3,248,604	17

Figure 3-5: China topped the list for web application attacks in Asia Pacific, with India and Japan in the second and third spots, respectively

3.3 / Top 10 Target Countries / The U.S. continues to be the largest target of attack traffic, with Brazil in second place for the second quarter in a row and the United Kingdom rounding out the top three. Attacks targeting the U.S. were down 9%, while Brazil saw a nearly 46% increase in web application attacks against their properties and the U.K. a 30% gain in attacks. Both China and Canada have fallen from the top 10 list this quarter, replaced by Spain and Singapore, which have both been on this list in the past. While these swings appear significant, they are within the norms we generally see for such traffic.

Top 10 Target Countries for Web Application Attacks, Q1 2017

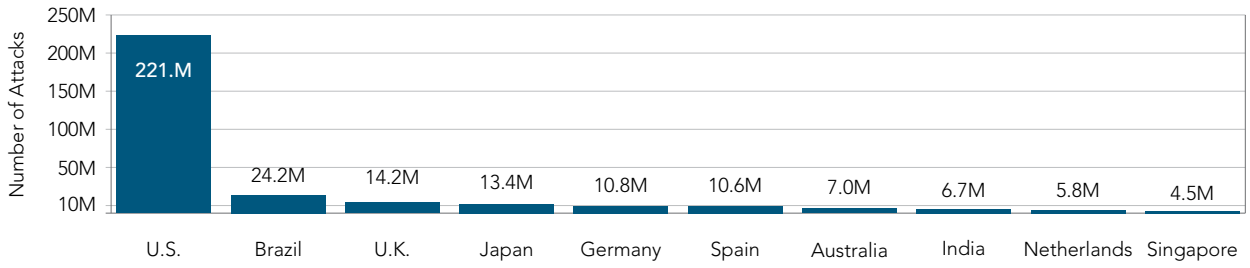

Figure 3-6: The vast majority of targets of web application attacks were in the U.S.

[SECTION]4
LOOKING FORWARD

The number of DDoS attacks has fallen in the last year, but have the risks been reduced as well? The answer is arguably no. If anything, the risks to the Internet as a whole and to targeted businesses in particular have both risen. Given the growth in capability of high-end attackers, the damage a sustained DDoS attack could cause increases daily. More and more often, it's not just the target that has to be concerned — other organizations may be affected by collateral damage from large DDoS attacks.

The size of the largest DDoS attacks jumped in 2016. Previously, the largest DDoS attacks were in the range of 100 Gbps, growing to 300 Gbps in first half of 2016, and finally into the 500-600 Gbps range in the third quarter, driven by Mirai. In addition to the attacks observed by Akamai, other organizations have seen DDoS attacks exceeding 1 Tbps. But the Mirai botnet is not only responsible for these large attacks — it's being used extensively in DDoS attacks of all sizes.

Attacks of this size easily overload the networks of their targets. In addition, they pose a problem for upstream networks that might not be able to handle the traffic, causing a multitude of organizations to be overwhelmed. It's like a crowded entry to a concert venue; a normal load might cause some headache, but the largest audiences not only overwhelm the venue, they also overflow into the roads and highways surrounding the area, affecting businesses and households for miles around. Instead of roads, it's the local loops and provider interconnects that are overwhelmed, unable to carry network traffic to organizations unlucky enough to be in the same region as the target.

Most botnets are not a single entity. For example, there are many Mirai-derived botnets using similar software, each a small fragment and distinct entity. There is constant fighting for control of the end nodes that comprise the botnets and the largest attacks are generally only seen when multiple distinct botnets target the same organization at once. One concern is that a unified command and control (C2) structure could emerge, either due to a new zero-day vulnerability or a takeover of the C2s of other similar botnets. Given the current capabilities of Mirai, such a super botnet could generate a DDoS attack of two Tbps in the near future. Additionally, Mirai's attacks are currently limited by the level of connectivity in their local networks. If these networks gain unfettered Internet access, the devices could be capable of emitting 20 times more attack traffic than we've seen to date.

The security community is taking measures to combat Mirai and other IoT-based botnets. As mentioned in the Emerging Trends section, Europol is helping coordinate global efforts to arrest the owners of the offending botnets. Some ISPs are taking measures to null route C2 traffic from botnets, dumping the bits before they leave the local network. Service providers and researchers are working to gain more insight into the structure of Mirai, in an attempt to limit its ability to spread and cause more damage.

It's short sighted to think of Mirai as the only threat, though. With the release of the source code, any aspect of Mirai could be incorporated into other botnets. Even without adding Mirai's capabilities, there is evidence that botnet families like BillGates, elknot, and XOR are mutating to take advantage of the changing landscape. In particular, the BillGates botnet family included the most recent Struts vulnerability[7] in its toolkit, very soon after the vulnerability was made public.

Finally, it's important to recognize that DDoS and the other threats from IoT are just one aspect of the threat landscape. Future *State of the Internet / Security* reports will examine traffic being sent to the APIs of web servers and explain how it could be an overlooked area of concern. Organizations may monitor the login page logs of their sites, but are they watching the traffic for their APIs? Site-to-site and business-to-business APIs may be a bigger target than most realize.

[SECTION]5 CLOUD SECURITY RESOURCES

5.1 / CLDAP DDoS Threat Advisory / On Oct. 14, 2016, the Akamai Security Operation Center (SOC) began mitigating attacks for what was suspected to be Connection-less Lightweight Directory Access Protocol (CLDAP) reflection. This new reflection and amplification method has since been confirmed by the Akamai SIRT and has been observed producing DDoS attacks, comparable to DNS reflection with most attacks exceeding 1 Gbps.

Similar to many other reflection and amplification attack vectors, CLDAP attacks would not be possible if proper ingress filtering was in place. Potential hosts are discovered using Internet scans. Filtering UDP destination port 389 can prevent CLDAP servers from being discovered and added to the attacks. Since October 2016, Akamai has detected and mitigated 50 CLDAP reflection attacks. Of those 50 attacks, 33 were single vector attacks using CLDAP reflection exclusively.

This ***advisory*** covers the distribution of these sources, methods of attack, and target industries observed.

Errata / Due to an error in the calculations for the maps and data for Figure 3-5: Web Application Attack Source Countries — Asia Pacific, Q4 2016 was missing data for Singapore. Singapore was the source of 1,644,483 attack events in Q4, which ranked it in fourth place for the Asia-Pacific region and 19th worldwide.

Endnotes /

[1] *https://www.extremetech.com/internet/247521-mirai-infamous-iot-botnet-now-forces-smart-appliances-mine-bitcoin*

[2] *https://arstechnica.com/security/2017/04/brickerbot-the-permanent-denial-of-service-botnet-is-back-with-a-vengeance/*

[3] *http://www.bankinfosecurity.com/ddos-for-hire-israel-arrests-two-suspects-a-9392*

[4] *https://www.grahamcluley.com/ddos-hire-arrests-europol-fbi/*

[5] *https://www.akamai.com/us/en/about/our-thinking/threat-advisories/connection-less-lightweight-directory-access-protocol-reflection-ddos-threat-advisory.jsp*

[6] *http://blog.erratasec.com/2017/04/mirai-bitcoin-and-numeracy.html#.WPoE3FPysSM*

[7] *https://blogs.akamai.com/2017/03/vulnerability-found-in-apache-struts.html*

***State of the Internet / Security* Team**
Martin McKeay, Senior Security Advocate, Senior Editor
Jose Arteaga, Akamai SIRT Lead, Data Wrangler — Attack Spotlight, DNS Water Torture
Amanda Fakhreddine, Editor
Dave Lewis, Senior Security Advocate — DDoS Activity, Web Application Attack Activity
Chad Seaman, Akamai SIRT — Attack Spotlight, DNS Water Torture
Wilber Mejia, Akamai SIRT — Attack Spotlight
Elad Shuster, Security Data Analyst, Threat Research Unit
Jon Thompson, Custom Analytics

Special thanks to Wendy Nather, Principal Security Strategist, Duo Security, for contributing to this quarter's report and to Jay Jacobs, Cyentia Institute, for his work on Figure 2-13: Attack Density and Trends

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

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*.
