



Industry Technical White Paper

July 17, 2017

ABSTRACT

On May 11, 2017 President Trump signed Executive Order 13800, Strengthening the Cybersecurity of Federal Networks and Critical Infrastructure, tasking the Department of Commerce and the Department of Homeland Security to lead an open and transparent process to identify ways to improve the resilience of the internet and communications ecosystem and reduce the threats perpetuated by botnets, particularly distributed denial of service attacks. In this technical white paper, the communications sector describes the botnet problem from the perspective of internet service providers (ISPs), identifies some challenges and opportunities, and then proposes several preliminary recommendations or actionable steps that ecosystem participants, including ISPs, should consider to mitigate the threats associated with botnets and automated attacks.

Communications Sector Coordinating Council

Table of Contents

Executive Summary	1
Internet Ecosystem and Communications Sector	3
Bots, Botnets and Associated Threats.....	7
Current Tools and Techniques	14
Emerging Solutions.....	18
Challenges and Opportunities	21
Industry Recommendations.....	29
Conclusion.....	31
Appendix A - Cyber Threat Reports.....	i
Appendix B – Threats from Botnets	iv
Glossary	vi

Executive Summary

A bot is a code used to seize control over a computer or a device to form a network of infected machines, known as a botnet. Many botnets are self-spreading and self-organizing networks of compromised machines that can be used to perform malicious activities in a coordinated way through command and control (C&C) channels. While bots are not new, the growing deployment of Internet of Things (IoT) devices amplifies their capability to create a large-scale global security threat.

In recognition of this growing global threat, on May 11, 2017, President Trump signed Executive Order 13800, *Strengthening the Cybersecurity of Federal Networks and Critical Infrastructure*,¹ tasking the Department of Commerce (DoC) and the Department of Homeland Security (DHS) to lead an open and transparent process to identify ways to improve the resilience of the internet and communications ecosystem and reduce the threats perpetuated by bots and botnets.

In this technical white paper, the communications sector, specifically internet service providers (ISPs) in this context, seeks to inform that process by describing the shared responsibilities of key participants in the internet ecosystem for mitigating the threats posed by botnets. It is a fallacy to believe that any single component of the internet ecosystem has the ability to mitigate the threat from botnets and other automated systems. While ISPs, as infrastructure owners and operators, play an important role in this ecosystem, so do the manufacturers of devices, developers of software, system integrators, edge providers, cloud service providers, and others. It will take the concerted effort of all members of this ecosystem to address fully the threats from bots and botnets.

The internet ecosystem has been working collaboratively to neutralize the threats from bots and botnets for years. In this paper, the Communications Sector Coordinating Council (CSCC) identifies a number of challenges of mitigating botnets, and opportunities for increased collaboration and cooperation among members of the internet ecosystem to address the problem including:

- Improving the efficiency of law enforcement process to take down botnets;

¹The White House Office of the Press Secretary, *Presidential Executive Order on Strengthening the Cybersecurity of Federal Networks and Critical Infrastructure* (May 11, 2017), available at <https://www.whitehouse.gov/the-press-office/2017/05/11/presidential-executive-order-strengthening-cybersecurity-federal>.

- Sharing of actionable cyber threat information;
- Reducing the dependency upon the use of network address translation (NAT) functions;
- Mitigating botnet traffic from foreign countries;
- Managing end-user notifications of malware infections;
- Defending against unsecured IoT devices;
- Combatting the use of fast flux domain name server (DNS) by botnets to hide their infrastructure; and
- Coordinating network-to-network network management.

As part of DoC and DHS's open and transparent process, the CSCC also proposes the following preliminary recommendations or actionable steps that ecosystem participants, including ISPs, should consider to mitigate the threats associated with bots, botnets, and automated attacks:

- Streamline the law enforcement process to take down botnets;
- Encourage continued migration to IPv6;
- Ensure that shared cyber threat information is actionable and tailored to meet recipients' needs;
- Network operators and end-users should include pre-negotiated provisions for traffic filtering in transit and peering agreements;
- Encourage the Internet Corporation for Assigned Names and Numbers (ICANN), registries, and registrars to adopt the fast flux mitigation techniques recommended by the Security and Stability Advisory Committee (SSAC);
- Improve botnet detection by encouraging the adoption and use of machine learning techniques;
- Ensure all end-points including IoT devices adhere to industry developed security standards;
- Ensure end-points are running up-to-date software; and
- IoT devices should use network isolation and/or network based filtering techniques for any communications to cloud-based services.

Internet Ecosystem and Communications Sector

The ecosystem supporting the internet, including the members of the communications sector providing internet access services is complex, diverse, and inter-dependent. To fully understand the threats that botnets pose, it is important to understand the ecosystem and stakeholders' relationships. This section provides a summary of the internet ecosystem and explains how the communications sector fits into the broader internet ecosystem in protecting critical infrastructure from threats from bots and botnets.

Internet Ecosystem

The internet ecosystem is a diverse, highly integrated system comprised of many stakeholders. The Internet Society (ISOC) describes the broad internet ecosystem as being made up of six primary communities as shown below.²

² Internet Society, *Who Makes the Internet Work: The Internet Ecosystem* (Feb. 3, 2014), available at <http://www.internetsociety.org/who-makes-internet-work-internet-ecosystem> (accessed July 16, 2017).

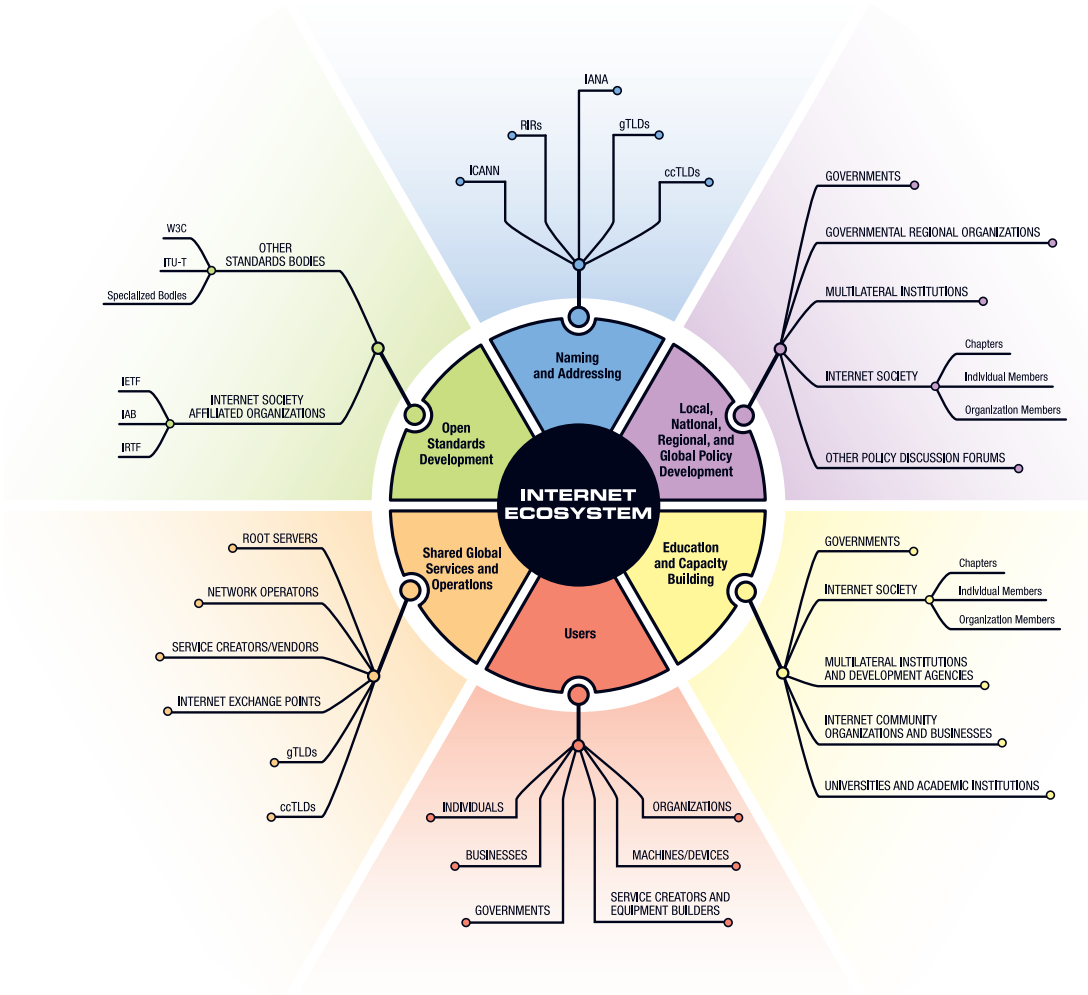


Figure 1 Internet Ecosystem

Source: Internet Society

The network operators, which are part of the communications sector, provide the “Shared Global Services and Operations” shown in Figure 1. When viewed solely from the network perspective, the internet ecosystem looks more like Figure 2.




Figure 2 Network View of Internet Ecosystem

In this context, the internet ecosystem is comprised of many machines/devices (e.g., smartphones, desktop computers, IoT devices, etc.) that connect to network service providers. The network service providers use a combination of transit and peering³ to provide internet connectivity to service creators (e.g., hosting, ecommerce, social media, enterprises, etc.). Many of the service creators are cloud-based, meaning that they operate a network of machines working together to provide a service. All of the parts work together to provide what is commonly referred to as the internet.

Communications Sector

Owners and operators of communications infrastructure (broadcast, cable, satellite, wireless, and wireline) comprise the communications sector. The communications sector is one of the 16 Critical Infrastructure/Key Resource (CI/KR) sectors identified in the DHS National Infrastructure Protection Plan (NIPP). This sector includes the network operators that provide internet access services. As part of a public/private partnership with DHS, the communications sector utilizes the Communications Sector Coordination Council (CSCC) and the Communications Information

³ Note: There is a glossary in Appendix B that provides more information on the technical terms used in this document.

Sharing and Analysis Center (Comm-ISAC) to help secure the communications networks CI/KR from harm.

The communications sector has a long history of cooperation within its membership and with federal government with respect to national security and emergency preparedness. This history distinguishes the communications sector from most other critical sectors identified in the National Infrastructure Protection Plan (NIPP). The sector exemplifies cooperation and trusted relationships that have resulted in the delivery of critical services when emergencies and disasters occur. This strong bond exists largely because of three organizations that have been created in response to earlier threats to the nation's critical infrastructure.

Policy - The National Security Telecommunications Advisory Committee (NSTAC). The NSTAC (www.ncs.gov/nstac/nstacthtml) was created in 1982 by Executive Order 12382. It provides a highly successful example of how industry helps direct government decisions around national security and emergency preparedness communications (NS/EP). NSTAC is comprised of up to 30 chief executives from major telecommunications companies, network service providers, and information technology, finance, and aerospace companies. Through a deliberative process, they provide the President with recommendations intended to assure vital telecommunications links through any event or crisis, and to help the U.S. Government maintain a reliable, secure, and resilient national communications posture. Key areas of NSTAC focus include: strengthening national security; enhancing cybersecurity; maintaining the global communications infrastructure; assuring communications for disaster response; and addressing critical infrastructure interdependencies.

Planning - Communications Sector Coordinating Council (CSCC). The CSCC was chartered in 2005 in order to: help coordinate initiatives to improve the physical and cybersecurity of sector assets; ease the flow of information within the sector, across sectors and with designated Federal agencies; and address issues related to response and recovery following an incident or event. The more than 40 members of the CSCC broadly represent the sector and include cable providers, commercial and public broadcasters, information service providers, satellite providers, undersea cable providers, utility telecom providers, service integrators, equipment vendors, and wireless and wireline owners and operators and their respective trade associations.

Operations - National Coordinating Center for Telecommunications (NCC)

Communications Information Sharing and Analysis Center (Comm-ISAC). In 1982, federal government and telecommunications industry officials identified the need for a joint mechanism to coordinate the initiation and restoration of national security and emergency preparedness telecommunications services. In 1984, Executive Order 12472 created the NCC. This organization’s unique partnership between industry and government advances collaboration on operational issues on a 24 X 7 basis and coordinates NS/EP responses in times of crisis. Since 2000, the NCC’s Communications Information Sharing and Analysis Center (Comm-ISAC), comprised of 51 industry member companies, has facilitated the exchange of information among government and industry participants regarding vulnerabilities, threats, intrusions, and anomalies affecting the telecommunications infrastructure. Industry and government representatives meet weekly to share threat and incident information. During emergencies, industry and government representatives involved with the response efforts meet daily, or even more frequently.

Bots, Botnets, and Associated Threats

Bot – a program that is installed on a system in order to enable that system to automatically (or semi-automatically) perform a task or set of tasks typically under the command and control of a remote administrator (aka bot master or bot herder).⁴

Botnet – a network of internet-connected end-user computing devices infected with bot malware and are remotely controlled by third parties for nefarious purposes.⁵

Bots are not a new phenomenon. It is important to note that not all bots are bad, and not all botnets are used for nefarious purposes. There are some good bots in environments like gaming and Internet Relay Chat (IRC). However, for the purposes of this paper, all mentions of bots and botnets will assume they are malicious or potentially malicious in nature.

A “botnet” is a network of bots working together with the capability of acting on instructions generated remotely. A typical botnet may range from a few thousand bots to hundreds of

⁴ Federal Communications Commission (FCC), Communications Security Reliability and Interoperability Council (CSRIC) III, *U.S. Anti-Bot Code of Conduct (ABCs) for Internet Service Providers*, (Mar. 2012), available at <https://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC-III-WG7-Final-ReportFinal.pdf> (accessed June 20, 2017).

⁵ *Id.*

thousands or even millions of bots. Bots and botnets are highly customizable and can be programmed to do many things, including: theft of personal and other sensitive information, spam, email address harvesting, distributed denial of service (DDoS) attacks, key-logging, hosting illegal content, and click fraud. These types of cyber-attacks are described in greater detail later in this paper.

Early bots used IRC to communicate to their C&C servers. Over time, bots and botnets have grown more sophisticated. For instance, bots and botnets have been made more resilient by incorporating peer-to-peer (P2P) architectures and protocols; domain name generating algorithms; hypertext transfer protocol (HTTP) to specific uniform resource locators (URL) within legitimate websites; sophisticated, hierarchical C&C infrastructures; and encryption. Each of these improvements has made it more difficult to identify and isolate bad traffic from legitimate network traffic.

Historically, bots infected desktop computers and servers, resulting in eventual detection and removal using antivirus software. In contrast, IoT devices often do not have a user interface (UI); are designed to run autonomously; and are connected either directly or indirectly to the internet. These devices do not lend themselves well to some traditional security protections. They may connect to the internet without a firewall and are usually placed on the same local area network (LAN) segment as other higher value targets. They are unlikely to run anti-virus software. In addition, they may be considered a low security risk since they are low cost and only process seemingly innocuous data. However, IoT devices are actually enticing targets for exploitation, as the devices provide computing power that can be utilized by bad actors, typically unnoticed by the owners, and are often “install and forget” equipment.

Large networks of IoT devices can become compromised by bots when connected to high-speed internet connections, which can cause significant damage. The October 2016 Mirai botnet DDoS attack against DNS provider Dyn is one of the more recent examples. The Mirai botnet exploited weak security in many IoT devices by continuously scanning the internet, looking for more IoT devices that were protected by factory default or hardcoded user names and passwords.⁶ As the Mirai botnet discovered vulnerable IoT devices, it loaded its malware onto the devices and began communicating with the C&C servers awaiting instructions. The Mirai botnet then was

⁶ Symantec Security Response, *Mirai: what you need to know about the botnet behind recent major DDoS attacks*, Symantec Official Blog (Oct. 27, 2016), available at <https://www.symantec.com/connect/blogs/mirai-what-you-need-know-about-botnet-behind-recent-major-ddos-attacks> (accessed June 20, 2017).

used to launch a large-scale DDoS attack against Dyn by instructing each infected device to flood the Dyn DNS servers with a high-volume of packets using the DNS service destination port (user datagram protocol (UDP) port 53) as well as flooding authoritative servers with numerous requests for invalid domain names.⁷ The attack prevented a number of Dyn’s customers from being able to access domain names served by Dyn DNS during the attack.

The Dyn attack was not an isolated incident. The peak attack size increased dramatically in a short period of time, rising from 500 Gbps in 2015 to 800 Gbps in 2016.⁸ The KrebsSecurity site was also hit by an attack in September 2016, which reached 620 Gbps. In fact, the Mirai botnet and other IoT botnets were in existence for some time prior to these attacks and generally used for performing smaller DDoS attacks.

Botnet Threats

As described above, bots and botnets are highly customizable, and as a result, can be programmed to do many beneficial things on the internet. However, they are often and increasingly, used for nefarious activities such as the types of attacks listed below.

- DDoS attacks;
- Data theft;
- Illicit content distribution;
- Brute force password guessing;
- Processing theft;
- Click fraud;
- Email spam; and
- Unauthorized gateway.

The remainder of this section, however, will focus on DDoS attacks. Descriptions of the other types of attacks listed above can be found in Appendix B.

⁷ Scott Hamilton, *Dyn Analysis Summary Of Friday October 21 Attack*, Dyn Blog (Oct. 26, 2016), available at <http://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/> (accessed June 20, 2017).

⁸ Arbor Networks, *12th Annual Worldwide Infrastructure Security Report*, Arbor Networks Special Report Vol. XII (2016), at p. 21, available at https://pages.arbornetworks.com/rs/082-KNA-087/images/12th_Worldwide_Infrastructure_Security_Report.pdf (accessed June 30, 2017).

DDoS attacks – a highly prevalent form of attack perpetrated by botnets – illustrate some of the many challenges of preventing attacks, as well as of preventing bots from compromising end-points.

DDoS attacks can be broken into four main categories:⁹

- Volumetric;
- Application/resource;
- State exhaustion; and
- Control plane.

Volumetric DDoS attacks consist of hundreds to hundreds of thousands of bots flooding the victim with packets, resulting in denial of the service to others. The attacks can be direct, where the bots send the packets addressed directly to the victim either with their own source IP address or a spoofed source IP address. Indirect attacks leverage a technique known as a reflective amplification attack, in which bots spoof the source IP address to be that of the intended attack target.¹⁰ The bots then send request packets to other services such as DNS, Character Generator Protocol (chargen), or Simple Service Discovery Protocol (SSDP) to trick the services to send responses toward the victim. Indirect or reflection attacks are often crafted to cause the service to send a response that is much larger than the bot's initial request, resulting in an amplification attack. In some circumstances, the amplifications can be thousands of times greater than the bots' initial request packets.

Application attacks tend to be lower volume traffic attacks than volumetric attacks. They are characterized by bots sending legitimate-looking application-level requests to a system to consume resources (e.g., CPU, disk access, data base lookups, etc.) and overwhelm the system, thereby preventing others from accessing it.

State exhaustion attacks leverage the fact that devices like servers, firewalls, and intrusion detection systems have limited capabilities to track the state of concurrent transactions. The

⁹ FCC CSRIC IV, *Remediation of Server-Based DDoS Attacks Final Report*, (Sept. 2014), available at [https://transition.fcc.gov/pshs/advisory/csric4/CSRIC_IV_WG5_Remediation_of_Server-Based_DDoS_Attacks_Report_Final_\(pdf\)_V11.pdf](https://transition.fcc.gov/pshs/advisory/csric4/CSRIC_IV_WG5_Remediation_of_Server-Based_DDoS_Attacks_Report_Final_(pdf)_V11.pdf) (accessed June 20, 2017).

¹⁰ Messaging, Malware and Mobile Anti-Abuse Working Group, *M3AAWG Introduction to Reflective DDoS Attacks* (May 2017), available at <https://www.m3aawg.org/sites/default/files/m3aawg-reflective-ddos-attack-intro.pdf> (accessed June 20, 2017).

bots leverage this limitation and consume all the state capabilities by opening many connections and not fully continuing those connections to completion.

Control plane attacks leverage the limitations of the internet protocols such as the Border Gateway Protocol¹¹ (BGP), IPv6,¹² and DNS protocol.¹³

A challenge with all types of DDoS attacks -- especially for ISPs -- is identifying them. Cyber criminals are rapidly devising more sophisticated botnets, making it harder to distinguish bad traffic from good traffic. The earliest forms of bots often transmitted their messages in clear-text, on well-known ports, to hard-coded IP addresses, thereby making the traffic both easy to identify and to block. Increasingly bots masquerade their traffic as application--level traffic (e.g., they make it look like regular web traffic or encrypted web traffic, use peer-to-peer techniques to avoid a single point of failure, or use VPNs to encrypt and tunnel their traffic to evade detection).

The Mirai botnet attack also leveraged the fact that there are millions of IoT devices all over the globe, and the attack traffic was generated from the far corners of the internet, sourced at the victims' locations. Level 3 Threat Research Labs reported that it observed over a million IoT devices participating in botnet attacks, and a large percentage were located in Taiwan, Brazil, and Columbia.¹⁴ The challenge for an ISP in detecting and blocking this traffic is that it does not originate on the ISP's network and may only transit a portion of the network, if it transits it at all. And even if there are bots on the network originating traffic, the volume of traffic from the bots may not be high enough to detect on the network.

Botnet attack traffic may look entirely normal. Much of it is reflective amplified attacks (which offer the best bang for the buck), frequently using well known common services such as DNS, network time protocol (NTP), and HTTP.

¹¹ K. Butler, et al, *A Survey of BGP Security Issues and Solutions*, Proceedings of the IEEE 98, no. 1 (Jan. 2010), at p. 100-122 (doi:10.1109/jproc.2009.2034031).

¹² Cisco, *IPv6 Extension Headers Review and Considerations [IP Version 6 (IPv6)]*, (Oct. 10, 2006), available at http://www.cisco.com/en/US/technologies/tk648/tk872/technologies_white_paper0900aecd8054d37d.html (accessed June 30, 2017).

¹³ Suranjith Ariyapperuma, and Chris Mitchell, *Security vulnerabilities in DNS and DNSSec*, Proceedings of Proceedings of The Second International Conference on Availability, Reliability and Security, ARES 2007, The International Dependability Conference - Bridging Theory and Practice, Austria, Vienna, available at <http://web.mit.edu/6.033/www/papers/dnssec.pdf> (accessed June 30, 2017).

¹⁴ Level 3 Research Labs, *Attack of Things!*, available at <http://www.netformation.com/level-3-pov/attack-of-things-2> (accessed June 20, 2017).

There are hundreds of different types of attacks within the five DDoS attack categories. Mirai itself has about a dozen DDoS attacks programmed into it. The botnet spread by scanning for open telnet ports (transmission control protocol port 23). Telnet is a clear text protocol and is extremely insecure and should not be used over the internet, but this is exactly how Mirai was spread. During the Dyn DNS attack, Mirai used DNS “water torture,”¹⁵ which it proxied through several well-known open resolvers (Google 8.8.8.8, for example). The attack on the KrebsOnSecurity¹⁶ website was designed to appear like the generic routing encapsulation (GRE) protocol.¹⁷ Both attacks could have been blocked by upstream internet transit providers. In the case of the Dyn attack, network service providers and the Comm-ISAC reached out to Dyn to offer assistance.

The KrebsOnSecurity attack being GRE-based could have been blocked by most ISPs. The Dyn traffic was proxied by well-known open resolvers, so rate limiting that traffic towards the Dyn IPs could have mitigated most of the effects of that attack. Brobot, which affected many U.S. financial systems, used HTTP and HTTPS for most of its attacks. Blocking it would require content examination and filtering, something ISPs generally do not do and cannot do for HTTPS without holding the end-user’s private keys. Malicious traffic that is encrypted (e.g., HTTPS) cannot be filtered.

The latest attacks illustrate the sophistication and scale that botnets have achieved. Botnets are detectable; the challenge is stopping them. The best way to stop them is to prevent their spread in the first place. The real challenge for the internet ecosystem in dealing with botnet threats is the remediation of infected end-points. Without either remediating the end-point or disconnecting the infected end-point from the internet, the threat from the infected end-point remains. Ensuring that end-points are running the latest software with the latest security patches is a recognized best practice for mitigating the spread of and threats from malicious and nefarious bots.

¹⁵ DNS water torture is an attack type where many end-points send queries for a victim’s domain with a random string prepended to the domain that overwhelms the victim’s authoritative DNS server and making the victim’s domain inaccessible.

¹⁶ See, <https://krebsonsecurity.com>.

¹⁷ KrebsOnSecurity, *KrebsOnSecurity Hit With Record DDoS* (Sept. 21, 2016), available at <http://krebsonsecurity.com/tag/gre-ddos/> (accessed July 16, 2017).

Most Botnet Traffic Originates Outside the United States

The threat landscape from botnets continues to evolve. According to threat intelligence companies, notable trends identified in the threat landscape in 2016 are that: 1) insecure IoT devices are a big source of DDoS attack traffic;¹⁸ and 2) the vast majority of the attack traffic originates from outside the United States.¹⁹

In 2016, attacks from IoT devices made headlines with the Mirai botnet attacks from improperly secured security cameras and their closed-circuit TV (CCTV) recorders (DVRs). As noted by Level 3 Threat Research Labs, many of the insecure cameras and DVRs were located in Taiwan, Brazil, and Columbia.²⁰ Shodan,²¹ a search engine that lets the user find specific types of IoT and other devices that are connected and visible on the public internet, reports (as of July 2017) 300K+ susceptible Hikvision devices connected directly to the internet, with the vast majority of those devices located in Brazil (45,000), India (36,000), China (34,000), Mexico (25,000), and South Korea (20,000).²²

While attributing the exact source of botnet attacks is difficult, it is almost always possible to determine the source country of the traffic. Numerous reports²³ indicate that the leading sources of attack traffic are China and other countries in Southeast Asia (e.g., Vietnam, Taiwan, and Thailand).²⁴

This is consistent with an earlier study that showed a strong correlation between devices used for botnet attacks and the country in which the devices reside. Such devices are typically running software without the latest security patches.²⁵ In one study, researchers analyzed six

¹⁸ Akamai, *State of the Internet Security Q4 2016 Report* (Winter 2016), available at <https://www.akamai.com/us/en/multimedia/documents/state-of-the-internet/q4-2016-state-of-the-internet-security-report.pdf> (accessed June 20, 2017).

¹⁹ Incapsula.com, *Global DDoS Threat Landscape Q1 2017* (Spring 2017), available at <https://www.incapsula.com/ddos-report/ddos-report-q1-2017.html> (accessed June 20, 2017).

²⁰ Level 3 Research Labs, *Attack of Things!*, available at <http://www.netformation.com/level-3-pov/attack-of-things-2> (accessed June 20, 2017).

²¹ See shodan.io (Shodan scans the internet indexing devices that respond to port scans on port 80, 8080, 443, 8443, 21, 22,23,161, 5060, 554, and other well-known ports).

²² Shodan, Search of "Hikvision," available at <https://www.shodan.io/search?query=hikvision> (accessed June 20, 2017).

²³ See Appendix A of this paper for data from different threat reports.

²⁴ Incapsula.com, *Global DDoS Threat Landscape Q1 2017* (Spring 2017), available at <https://www.incapsula.com/ddos-report/ddos-report-q1-2017.html> (accessed June 20, 2017).

²⁵ Hadi Asghari, Michael Ciere, and Michael J.G. Van Eten, *Post-Mortem of a Zombie: Conficker Cleanup After Six Years*, In USENIX The Advanced Computing Systems Association, Proceedings of 24th USENIX Security Symposium, Washington, D.C. (Aug. 2015), available at <https://www.usenix.org/system/files/conference/usenixsecurity15/sec15-paper-asghari.pdf> (accessed June 20, 2017).

years of longitudinal data from the sink-hole of Conficker, one of the largest botnets ever seen, to assess the impact on botnet mitigation of national anti-botnet initiatives, aimed at getting end-users to clean infected end-user machines. They found that peak infection levels strongly correlate with software piracy. This implies that countries with a higher number of end-users running unlicensed copies of software tend to have higher numbers of bots because those assets have a lower percentage of registered users getting security patches.

A similar pattern was seen with the Mirai botnet, which exploited the fact that a class of IoT devices shipped with well-known, default login credentials that end-users rarely change. Vulnerabilities with at least one of the manufacturers were reported as far back as 2013.²⁶ Only after the Mirai botnet attack was reported did the manufacturer in question provide a firmware update to address the vulnerabilities, and, even then, it required manual intervention by device end-users to update the firmware, as the devices did not support an automated manner for securely updating their software.

Current Tools and Techniques

Application of Cybersecurity Framework against Botnets

The Cybersecurity Framework, developed by National Institute of Standards & Technology (NIST),²⁷ is a voluntary risk-based “set of industry standards and best practices to help organizations manage cybersecurity risks.” The Framework is composed of five functional areas – 1) Identify, 2) Detect, 3) Protect, 4) Respond, and 5) Recover. The leading ISPs use the Framework as part of their overall cyber risk management processes to address the threats posed by bots and botnets against their networks.

Identify

In the Framework, the first step is **identifying** both what needs to be protected and what are the cyber threats. The Federal Communications Commission’s (FCC) Communications Security,

²⁶ Department of Homeland (DHS) Security Office of Cybersecurity and Communications, *Vulnerability Note VU#800094 - Dahua Security DVRs contain multiple vulnerabilities* (Dec. 4, 2013), available at <http://www.kb.cert.org/vuls/id/800094> (accessed June 20, 2017).

²⁷ National Institute of Standards and Technology, *Cybersecurity Framework* (May 25, 2017), available at <https://www.nist.gov/cyberframework> (accessed June 20, 2017).

Reliability and Interoperability Council (CSRIC) IV Working Group 4 final report, *Cybersecurity Risk Management and Best Practices*, provides implementation guidance on the use of the Framework for network service providers. ISPs, as part of the critical infrastructure, have identified that they need to protect their core networks from cybersecurity threats such as bots and botnets. ISPs may also, as part of a managed security service, protect their customers from the harms of cyber threats.

In addition to identifying what needs to be protected, network service providers use the Framework and other tools to identify the threats. The first step is identifying the attack surfaces of the assets to be protected and then identifying the known attack vectors. This information is continuously synthesized with threat intelligence data to ensure comprehensive coverage and to identify, and ultimately address, new vulnerabilities. Obtaining high-quality cyber threat data is one of the most important aspects of implementing and running a strong botnet mitigation program. For the program to be effective, near zero false positive data is needed. False positives can greatly increase a network service provider's operating costs, impact its customer satisfaction, and damage its brand. As outlined in the CSRIC V Working Group 5 report on *Cybersecurity Information Sharing*,²⁸ network service providers have developed an information sharing ecosystem to both use and share cyber threat indicator information from an array of sources, to identify botnets and their associated threats. Included in this ecosystem are trusted third-party (TTP) data feeds, information from DHS including its Automated Information Sharing (AIS) system, and inter-sector information sharing.

Detect

As outlined in the Framework, **detection** of threats and attacks is the next step in protecting networks from botnet attacks. As described earlier, botnet attacks come in many forms, so detecting them requires an array of tools and techniques tailored for each kind of attack. Regardless of the type of botnet attack, network service providers use a core set of techniques, including packet sampling, signature analysis, and heuristic or behavioral analysis.

Many botnets attempt to disguise their traffic as normal internet traffic. This makes it particularly difficult to detect highly distributed botnets or low-volume traffic botnets, as the

²⁸ FCC CSRIC V, *Working Group 5: Cybersecurity Information Sharing*, Final Report (Mar. 15, 2017), available at <https://www.fcc.gov/files/csric5-wg5-finalreport031517pdf> (accessed June 20, 2017).

traffic will be below the alarm thresholds on any single operator's network. For example, during the Mirai Dyn DNS waterboarding attack, the attackers proxied their requests through well-known open DNS resolvers.²⁹

Protect

Network service providers use a variety of techniques to **protect** their networks from attacks and undertake measures to help their customers protect themselves from attacks.

Network service providers use different filtering techniques to directly protect their network infrastructure (e.g., routers, servers). Bots often spoof the source IP address in the attack packets. This is typically seen in network reflection attacks, but as seen in high volume attacks such as the Mirai botnet or Dyn attack, this can be accomplished even without IP spoofing. Regardless, as a best common practice, most, if not all, network service providers perform network filtering for spoofed IP addresses.³⁰

Network service providers also use a combination of other filtering techniques such as Access Control Lists (ACLs), traffic policing, black holing, and sink holing in their networks to filter known botnet traffic. These techniques can be effective for neutralizing the C&C traffic for client-server botnets. This is less effective against more advanced botnets that use peer-to-peer architecture, encryption, and/or fast flux DNS techniques for their C&C channel. Fast flux is a DNS technique used by botnets to hide phishing and malware delivery sites behind an ever-changing network of compromised hosts acting as proxies.

Network service providers also have made large investments in DDoS scrubbing systems to "scrub" out DDoS attacks against their networks and their customers who have purchased DDoS mitigation services. DDoS scrubbing systems rely upon diverting *the victim's* traffic through the scrubber "on-demand" to filter out attack traffic from good traffic, and then place it back on the provider's network to send it to its original destination. Network service providers use a combination of in-house scrubbing systems and third-party scrubbing systems via contracts with

²⁹ Scott Hamilton, *Dyn Analysis Summary Of Friday October 21 Attack*, Dyn Blog (Oct. 26, 2016), available at <http://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/> (accessed June 20, 2017).

³⁰ P. Ferguson and D. Senie, *Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing*, Best Current Practice (BCP) 38 (May 2000), available at <https://tools.ietf.org/html/bcp38> (accessed June 20, 2017); F. Baker, and P. Savola, *Ingress Filtering for Multihomed Networks*, BCP 84 (Mar. 2004), available at <https://tools.ietf.org/html/bcp84> (accessed June 20, 2017); and Mutually Agreed Norms for Routing Security (MANRS), *Participants* (Mar. 6, 2015), available at <https://www.routingmanifesto.org/participants/> (accessed June 20, 2017).

third party DDoS mitigation providers. However, network service providers do *not* have the capacity to scrub all traffic all of the time.

In addition to scrubbing traffic, many providers use the Flowspec³¹ capabilities of BGP to dynamically block easily identifiable traffic on the router. The traffic is usually blocked using the basic five-tuple of values found in IPFIX³² (source and destination IP, source and destination port, and protocol). Flowspec is advantageous in that BGP updates can be made and withdrawn fairly quickly in the network, allowing for faster mitigation.

Network service providers also can provide specific tools and services to their customers to protect themselves, including end-point anti-virus software and home gateways with integrated security.³³ Large ISP customers operating stub networks or edge providers also can use a technique to mitigate DDoS attacks known as Anycast, which allows multiple hosts or end-points to have the same IP address. By geographically distributing these hosts, the magnitude of the DDoS attack needs to be significantly larger to account for the distributed hosts and succeed at disrupting the site or service. Anycast services can be deployed by edge providers or purchased from DDoS mitigation partners.

Several network service providers also offer a suite of managed security services including but not limited to the DDoS scrubbing services mentioned above. These can include capabilities such as network based firewalls, mobile device management services, threat analysis and event detection, secure VPN connectivity to the cloud, and web and email security.

Respond & Recover

Today, as outlined in the Cybersecurity Framework, when a network service provider detects malicious traffic from a bot either on its network or toward an end-point on its network, it **responds and recovers** as necessary. The response consists of mitigating the impact from the malicious traffic, and, if necessary, remediating the infected end-point.

To mitigate the malicious traffic, the network service provider must first determine the scope of the impact from the malicious traffic. For malicious traffic that is impacting its network or its

³¹ Leonardo Serodio, *Traffic Diversion Techniques for DDoS Mitigation using BGP Flowspec* (May 2013), available at <https://nanog.org/sites/default/files/wed.general.trafficdiversion.serodio.10.pdf> (accessed July 7, 2017).

³² B. Claise, B. Trammell, and P. Aitken, *Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information*, IETF Tools (Sept.2013), available at <https://tools.ietf.org/html/rfc7011> (accessed July 7, 2017).

³³ McAfee, *McAfee Web Gateway*, available at <https://www.mcafee.com/us/products/web-gateway.aspx> (accessed July 7, 2017).

ability to deliver service, the network service provider will need to work to filter out the malicious traffic using one of the filtering techniques (e.g., ACL, black hole, sink hole, or scrub) described earlier. In addition, if the malicious traffic is inbound toward its network, the network service provider may contact the upstream network and ask it to filter the traffic emanating from that network.

For malicious traffic that is determined to be emanating from a customer end-point on its network, the network service provider, as recommended in the voluntary Anti-Bot Code of Conduct for Internet Service Providers (ABC for ISPs)³⁴ will:

- **Detect** – identify and detect botnet activity in the ISP’s network or on behalf of enterprise customers who have purchased services from the ISP to determine potential bot infections on end-user devices;
- **Notify** – notify end-users, including potentially both consumers and enterprise business clients of suspected bot infections;
- **Remediate** – provide information to end-users about how they can remediate bot infections and/or actively assist enterprise business clients in remediating the impacts of botnets; and
- **Collaborate** – provide feedback and experiences learned to other ISPs.

Emerging Solutions

The internet ecosystem is continuing to improve its ability to mitigate the attacks from botnets. Efforts are underway to improve both detection and mitigation capabilities.

Technological Approaches. A large number of malware uses a technique known as a domain generation algorithm (DGA) to periodically generate a large number of domain names that can be used as rendezvous points for their C&C servers in an attempt to obfuscate the botnet’s true infrastructure. Currently, security investigators can work to reverse engineer the DGA used by each variant of malware. The reverse engineering can be a time-consuming process, and is often an ineffective whack-a-mole approach. To address this issue, industry has been investigating how to apply machine learning to automate the process and work in real-time as

³⁴ Messaging Malware Mobile Anti-Abuse Working Group (M3AAWG), *ABCs for ISPs*, available at <https://www.m3aawg.org/abcs-for-ISP-code> (accessed June 20, 2017).

the malware registers domain names with an internet registry. Efforts are underway to commercialize and integrate machine learning for botnet detection into network protection products.

Newer botnets now often use encryption (e.g., TLS³⁵) to hide their C&C channel. The Secure Sockets Layer SSL Blacklist (SSBL) project³⁶ illustrates that even though the botnet is using encryption, it is still possible to detect the botnet. It is possible to identify the bot's C&C traffic by inspecting the malicious SSL certificates to generate a unique SHA-1³⁷ fingerprint for each botnet using deep packet inspection (DPI). Efforts are underway to commercialize this approach and integrate the methods into network protection systems to allow for real-time fingerprinting and mitigation of botnets.

In addition, researchers are developing the use of tarpits at network scale to slow the propagation of botnets.³⁸ Researchers are investigating how to turn unused IP address space into botnet tarpits.³⁹ The basic idea is to route all inbound traffic that is addressed to the unused IP addresses to the tarpit. The tarpit has a set of programmed rules for how to respond, and thereby extends the time it takes for a botnet to work its way up the kill chain.⁴⁰ By extending the time, the targets of the attack have more time to determine what additional defensive measures need to be put in place to neutralize the attack, if any.

In addition to tarpits, network providers have undertaken efforts to determine how to leverage the features of Software Defined Networks (SDNs) to help mitigate attacks from botnets. SDNs provide the capability to dynamically create overlay networks. When combined with other network partitioning techniques and technology, it becomes possible to dynamically create virtual lanes for the different IP-based services. With this approach, IoT providers can work with network service providers to create end-to-end virtual lanes from the IoT device through the network to the cloud-based service. This process ensures a compromised IoT device cannot

³⁵E. Rescorla and N. Modau, *Datagram Transport Layer Security Version 1.2*, IETF Tools (Jan. 2012), available at <https://tools.ietf.org/html/rfc6347> (accessed June 20, 2017).

³⁶SSL Blacklist, *SSL Blacklist*, available at <https://sslbl.abuse.ch/blacklist/> (accessed June 20, 2017).

³⁷SHA-1 – Secure Hash Algorithm 1 is a cryptographic hash function that generates a 20 byte hash key used by many security applications and protocols including TLS and SSL as part of encrypting data.

³⁸Labrea, *Tom Liston Talks about Labrea*, available at <http://labrea.sourceforge.net/Intro-History.html> (accessed July 17, 2017).

³⁹Tarpits are defensive measures against attacks where the server purposely delays incoming connections to make spamming and broad scanning less effective.

⁴⁰Eric Hutchins, Michael Cloppert, and Rohan Amin, *Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains*, CND Papers (Nov. 21, 2010), available at <http://papers.rohanamin.com/?p=15> (accessed July 7, 2017).

communicate with unauthorized endpoints. In other words, a compromised device could not be used in a DDoS attack or send information to non-authorized hosts. The Network Slicing feature in 5G networks is a good example of this,⁴¹ and similar approaches are being investigated for SDN-enabled wireline networks.

Collaboration Initiatives. Several industry-led initiatives are underway to improve automated cyber threat information sharing. The Cybersecurity Information Sharing Act (CISA), enacted in 2015, and the subsequent rollout of the DHS Automated Information Sharing (AIS) capability are helping to facilitate machine-to-machine (M2M) initiatives.

There are at least two other automated M2M sharing initiatives that may be useful in combatting botnets. Both have a common goal of ensuring that the cyber threat information being shared is “actionable” by the recipient. The paradigm in the past often has been for networks to try to build better protection at their network ingress points. These initiatives share information with neighboring networks to mitigate the threat as close to the source of the malicious traffic as possible.

The Internet Engineering Task Force (IETF) is developing a protocol called DDOS Open Threat Signaling (DOTS)⁴² for the real-time exchange of DDoS-related telemetry between DDoS mitigation network elements. The IETF DOTS protocol is working to improve the cooperation between DDoS attack victims and parties that can help in mitigating such attacks. The protocol will support requests for DDoS mitigation services and status updates across inter-organizational administrative boundaries (e.g., network-to-network).

The Messaging Malware Mobile Anti-Abuse Working Group (M3AAWG) DDoS special interest group⁴³ members are collaborating on a similar endeavor. M3AAWG is developing an application program interface (API), data store, and open source reference implementations for network service providers to share DDoS threat indicators for the purpose of identifying sources of DDoS attack traffic, but not for mitigating attacks in real time. M3AAWG’s approach allows network service providers to share the source IP addresses for the inbound IP flows that their DDoS detection systems identify in an anonymous fashion with the network on which the DDoS

⁴¹ See 5G Americas, *Network Slicing for 5G Networks & Services*, available at http://www.5gamericas.org/files/3214/7975/0104/5G_Americas_Network_Slicing_11.21_Final.pdf (accessed July 7, 2017).

⁴² IETF, *DDoS Open Threat Signaling (dots)*, available at <https://datatracker.ietf.org/wg/dots/about/> (accessed June 20, 2017).

⁴³ M3AAWG, *M3AAWG Issues New Papers Explaining Password Security, Multifactor Authentication, Encryption Use and DDoS Safeguards; Announces Leadership and Committee Chairs*, Press Release (Apr. 4, 2017), available at <https://www.m3aawg.org/news/rel-leadership-papers-2017-04> (accessed June 20, 2017).

attack originated. This allows network operators to clean up the sources of DDoS attack traffic. By sharing only the source IP address, this approach is compatible with most of the global privacy laws with respect to the sharing of identifiable information.

Challenges and Opportunities

Cybersecurity is shared responsibility. Reducing the threats from bots, botnets, and their automated attacks requires the cooperation and collaboration by all members of the internet ecosystem. This section identifies a number of areas where the threats presented by bots and botnets can be reduced with better cooperation and collaboration by members of the internet ecosystem.

Botnet Takedowns

Challenge – No technique is more effective than law enforcement actions that lead to the arrest of the perpetrators. This is the only solution that addresses the root cause of the problem, and not just a symptom. Unfortunately, executing a botnet takedown requires significant upfront forensic analysis and careful coordination among many stakeholders, often across international borders. A limiting factor in the overall velocity of botnet takedowns is the lack of law enforcement resources. The other challenge is that most botnets are international in nature, requiring resource-intensive and time-consuming cooperation between nations.

Opportunity – Additional law enforcement resources and streamlining international processes would aid the overall botnet takedown process.

Actionable Cyber Threat Information

Challenge - Network service providers must have both accurate and actionable cyber threat information to be able to quickly neutralize botnets. For information to be actionable, the cyber threat indicator has to be correlated to a single end-point. Many of the data feeds used and shared by enterprise are long-term IP reputation lists of little value to network service providers that operate networks with a large set of subscribers that have dynamically assigned IP addresses with short leases. This means the cyber threat indicator must be timely and either include the current IP address or the IP address and a time-stamp of the malicious activity.

The same is true for IP addresses of the botnet C&C servers. C&C servers often do *not* have a static IP address. Often the C&C servers are on shared hosts where a single IP address is shared by multiple hosts. In addition, the C&C servers may have a pool of IP addresses or shared hosts that they rotate through.

Network service providers need a single, highly reliable, near-term indication that an IP address has generated malicious traffic or has been scanned to show exposed vulnerable services, as well as the compromised hosts.

Opportunities - Experts agree that cyber threat information needs to be timely and targeted to be effective. The cyber information sharing initiatives of the IETF's DOTS Working Group and the M3AAWG DDoS SIG are steps in the right direction. DHS's AIS⁴⁴ also provides an opportunity to improve and enhance the timely and tailored sharing of cyber threat indicators to meet recipients' needs.

Network Address Translation

Challenge – Wireline ISPs operating IPv4 networks typically provide a residential subscriber with a single public IPv4 address. The residential subscriber often uses a home router that includes a network address translation (NAT) function, which allows them to share their one public IPv4 address with multiple devices in the home.

When an ISP receives information about a residential subscriber sending malicious traffic, that information, at best, can only contain the IPv4 address assigned to the customer and not that of the actual end-point behind the home router. The use of NAT technology makes it difficult for the ISP to identify the specific device in the subscriber's home that is sending malicious traffic.

Opportunity - IPv6 eliminates the need to use NAT for IP address sharing, as every device connected to the internet can have a publicly routable IPv6 address. While not a panacea, the elimination of NAT routers may make it easier to identify end-devices transmitting malicious traffic under certain circumstances, and to filter the suspect traffic appropriately. As of June

⁴⁴ DHS, *Automated Indicator Sharing (AIS)*, available at <https://www.dhs.gov/ais> (accessed June 20, 2017).

2017, IPv6 adoption by network providers was approximately 19% globally,⁴⁵ and 35% and growing within the U.S.

Off-Net Traffic

Challenges - As widely distributed global networks, most bots and their C&C servers are outside the network service provider's network and administrative control. In fact, numerous reports make clear that the overwhelming majority of botnet traffic originates outside the U.S.⁴⁶

Furthermore, in most cases, only a small portion of a network service provider's end-points may be infected by any single botnet, and the amount of traffic generated by the botnet on the network will be miniscule. This small amount of traffic can be very difficult to detect as it will not trigger many of the network monitoring thresholds that a network service provider has in place.

Opportunity - To address both of these challenges requires collaboration among network service providers, as one of the most effective measures is to filter the traffic as close to the device infected with the bot. Any transit or peering agreements should include language that addresses availability and scrubbing of traffic to allow for network operators to ask the upstream provider(s) to filter malicious traffic.

End-User Notifications

Challenge - Notifying and getting end-users to take action continues to be a challenge. There are multiple ways that members of the internet ecosystem can notify an end-user:⁴⁷

- Email;
- Telephone call;
- Postal mail;

⁴⁵ Google, *IPv6 Adoption* (June 18, 2017), available at <https://www.google.com/intl/en/ipv6/statistics.html#tab=per-country-ipv6-adoption&tab=per-country-ipv6-adoption> (accessed June 20, 2017).

⁴⁶ Incapsula.com, *Global DDoS Threat Landscape Q4 2016* (Winter 2017), available at <https://www.incapsula.com/ddos-report/ddos-report-q4-2016.html> (accessed June 20, 2017).

⁴⁷ Michael Glenn, *Malware Notification and Remediation Tools and Techniques*, CenturyLink presentation to NIST Workshop: Technical Aspects of Botnet (May 30, 2012), available at https://www.nist.gov/sites/default/files/documents/itl/csd/centurylink_malware_notification_and_remediation.pdf (accessed June 20, 2017).

- Text message;
- Web browser notification;
- Walled garden; and
- Other Methods.⁴⁸

A study commissioned by M3AAWG to determine the effectiveness of various notification and remediation methods showed that the two most effective methods are a telephone call to the device user and postal mail.⁴⁹ The growing use of IoT devices in homes presents new challenges in notifying end-users. IoT devices often have limited user interfaces, thus negating a number of the notification methods (web browser, walled garden, etc.). This is further compounded by the fact that an ISP can only notify an end-user that “a device” in their home is infected, and cannot identify the specific corrupted device.

Opportunities – Various measures exist to improve device identification going forward. Better designed IoT devices that adhere to industry standards such as those being developed by the Open Connectivity Foundation (OCF)⁵⁰ is one avenue to improve security. And, as noted earlier, network operator support for IPv6 will aid in both the identification of the infected device, as well as notifying the user of the device.

Fast Flux DNS

Challenge – The use of fast flux⁵¹ by malware and botnets to hide their infrastructure continues to grow. Fast flux is a DNS technique where numerous IP addresses associated with a single domain name are swapped in and out with extremely high frequency. Fast flux effectively hides the computers or servers that are performing the malicious attacks from being detected. Fast flux makes cutting off contact of the bots to the C&C servers difficult or impossible by IP address filtering alone.

Opportunity – In 2008, the ICANN Security and Stability Advisory Committee (SSAC) published a security advisory that made a number of mitigation recommendations to address fast flux DNS

⁴⁸ Other methods may include social media message, alert to the TV via the set-top-box, direct deposit voicemail message, etc.

⁴⁹ Georgia Tech Researchers, *DNS Changer Remediation Study*, Presentation to M3AAWG 27th General Meeting, San Francisco, CA (Feb. 19, 2013), available at https://www.m3aawg.org/sites/default/files/document/GeorgiaTech_DNSChanger_Study-2013-02-19.pdf (accessed June 20, 2017).

⁵⁰ See Open Connectivity Foundation, available at <https://openconnectivity.org/> (accessed June 20, 2017).

⁵¹ ICANN Security and Stability Advisory Committee (SSAC), *SAC 025 SSAC Advisory on Fast Flux Hosting and DNS* (Mar. 2008), available at <https://www.icann.org/en/system/files/files/sac-025-en.pdf> (accessed June 20, 2017).

techniques. Among its findings and recommendations, the SSAC encouraged ICANN, registries, and registrars to consider the fast flux mitigation practices in the advisory.

Since that time, advancements in machine learning have been applied to detecting botnets using fast flux DNS techniques. Advancements in the application of machine learning to detect botnets that are making changes to DNS entries enables automation and integration into botnet detection systems.

Insecure IoT Devices

Challenge – As discussed throughout this paper, the growing installed base of IoT devices is making such devices attractive targets for cyber criminals to infect with bot code. A good example is the recent Mirai botnet attack, in which unsecured, internet-connected IoT security cameras were infected to generate a massive DDoS attack. This is not a new phenomenon; the problem has been around for years, as for years, many consumer-grade home routers shipped with known vulnerabilities that have been exploited to generate large-scale DNS amplification attacks.

The types of known vulnerabilities⁵² found in many IoT devices on the market today include:

- Shipping IoT devices with out-of-date software containing known vulnerabilities and lacking the capability for an automated software update;
- Protection only by factory default or hardcoded user names and passwords;
- Unauthenticated communications;
- Unencrypted communications; and
- Lack of mutual authentication and authorization.

Insecure IoT devices present a unique challenge as once they are compromised it is often impossible for the end-user to detect that they have been compromised and, as noted earlier, it is difficult for a network service provider to notify the end-user that their device has been compromised. Even after the end-user is aware of the compromise, it is often impossible to

⁵²Broadband Internet Technical Advisory Group (BITAG), *Internet of Things Security and Privacy Recommendations* (Nov. 2016), available at [http://bitag.org/documents/BITAG_Report_-_Internet_of_Things_\(IoT\)_Security_and_Privacy_Recommendations.pdf](http://bitag.org/documents/BITAG_Report_-_Internet_of_Things_(IoT)_Security_and_Privacy_Recommendations.pdf) (accessed June 20, 2017).

remediate the problem due to either the lack of a software update and/or lack of automated software updates.

Opportunity - IoT devices can be better secured through the use of network/path isolation.⁵³ Network/path isolation techniques (VPNs, VLANs, policy based routing, network slicing, etc.) can be used to create independent logical traffic paths. These independent logical traffic paths ensure the IoT traffic can only reach the designated endpoints. This helps to mitigate the impacts of any malicious traffic that a compromised IoT device may send.

With the advances in network function virtualization (NFV) and SDNs, opportunities exist for IoT manufacturers to design devices to use network/path isolation techniques as part of their service. Additionally, opportunities exist for network service providers to offer network/path isolation as a service to IoT providers or end-users for their IoT devices.

Amplification Attacks

Challenge - An amplification attack is a type of DDoS attack that takes advantage of the fact that a small query such as a DNS query can generate a much larger response. When combined with source address spoofing, an attacker can direct a large volume of network traffic to a target system. The asymmetric nature of amplification attacks makes it the preferred choice for DDoS attacks. Amplification attacks often leverage UDP based protocols such as the DNS protocol, network time protocol (NTP), character generator (CharGEN), and quote of the day (QOTD). Approximately 15 internet protocols are susceptible to amplification attacks.⁵⁴ Internet engineers developed an extension to the DNS protocol, called DNS Security (DNSSEC) to address DNS vulnerability to DNS cache poisoning. Unfortunately, a side effect of this fix is that the security extension to DNS makes the DNS responses much larger and helps to further amplify the attack.

The implementation of source address validation (SAV)⁵⁵ as recommended in IETF BCP 38/84 prevents amplification attacks with spoofed source addresses. Although most large U.S.

⁵³Cisco, *Network Virtualization--Path Isolation Design Guide* (July 22, 2008), available at http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Network_Virtualization/PathIsol.html (accessed June 20, 2017).

⁵⁴United States Computer Emergency Readiness Team (US-CERT), *UDP-Based Amplification Attacks*, Alert (TA14-017A) (Nov. 4, 2016), available at <https://www.us-cert.gov/ncas/alerts/TA14-017A> (accessed June 20, 2017).

⁵⁵ SAV has been a best practice by ISPs for a long time (see IETF 2267 published in 1998), but due to the difficulty of implementing SAV in some commercial situations it may not be fully implemented across ISPs' networks.

network service providers⁵⁶ have implemented source address validation, approximately 30% of the overall IP address space is still spoofable.⁵⁷

Opportunity - The use of IP filtering or source address validation (SAV) as outlined in the IETF's best common practices (BCP) 38 and 84 for spoofed IP addresses is a proven technique to mitigate DDoS amplification attacks using spoofed source addresses.

The Mutually Agreed Norms for Routing Security (MANRS)⁵⁸ is an industry-led effort to codify a set of shared values for network operators into a set of definitions and ideal behaviors. MANRS recommends the implementation of anti-spoofing filtering to prevent packets with incorrect source IP addresses from entering or leaving the network. To date, over 45 network operators are participating in MANRS. The opportunity exists to get the spoofable address space to near zero with every network operator participating in MANRS.

Network-to-Network Coordinated Network Management

Challenge - Although network management may sound simple and desirable, it is not without challenges, especially given the negative impact on internet end-users. Ideally botnet mitigations would be fast and directed at the source of the attack. Advancements in how networks are architected using SDNs and the use of automated M2M sharing of cyber threat indicators start to make it technically viable for network operators to automate the coordination of their botnet mitigations and reduce the response time to when either a malicious bot is detected on a network or a botnet is initiating an attack. But there are challenges, ranging from technical to contractual, and policy issues.

The technical challenges include both detection and mitigation. Without a source of ground truth for what is and isn't botnet traffic, given botnet traffic is often designed to look like normal internet traffic, there is the potential for false positives. Even with a source of ground truth, botnet mitigation methods will vary from network to network due to inherent differences in

⁵⁶MANRS, *Participants* (Mar. 6, 2015), available at <https://www.routingmanifesto.org/participants/> (accessed June 20, 2017).

⁵⁷Center for Applied Internet Data Analysis, *State of IP Spoofing*, available at <https://spoofer.caida.org/summary.php> (accessed June 20, 2017).

⁵⁸MANRS, *Mutually Agreed Norms for Routing Security (MANRS) Document* (Sept. 8, 2016), available at <http://www.routingmanifesto.org/manrs/> (accessed June 20, 2017).

how the networks are designed and built, as well as the differences in service level agreements between network service providers and their customers.

Blindly mitigating botnets through the use automation is fraught with risks. There are many cases where a command and control server is not owned or completely under the control of the bot operator such as: 1) shared server DNS, 2) shared IPs, and 3) public websites.⁵⁹ Blindly applying a botnet mitigation method such as filtering the IP address would prevent all the services that share the resource (e.g., DNS, shared server, or service) from being accessible. The challenge is not limited to shared resources. Without full knowledge of the service level agreement in place between the network service provider and customer, a network service cannot blindly filter the traffic to that end-point.

In addition, within the telecom/ISP industry there is an emerging trend toward the adoption of SDN, which is still in its infancy, but generally describes the automation of management and orchestration of network assets and services. Typically, this includes the coupling of big data frameworks that leverage advanced analytics and machine learning to serve as feedback loops for these SDN-driven networks to predict, recommend, and prescribe in an effort to improve responsiveness and resilience of their assets and services. Such implementations vary widely in terms of capability and maturity across providers, and in most cases reflect highly protected intellectual property that provides a uniquely competitive experience and offerings. Nevertheless, such an ecosystem could be used as an attack mitigation strategy. Deployment of SDN and these tools is well beyond the conceptual stages; it is the complexity and cost of global implementation across highly heterogeneous networks that stand as obstacles to providers' speed in implementing them.

Opportunity –Better collaboration and coordination can reduce the time that it takes to respond to cyber threats. As mentioned earlier, industry is developing solutions such as the IETF DOTS, M3AAWG DDoS SIG's information sharing pilot, and an information sharing pilot being led by CTIA that will reduce the response time by sharing "actionable" cyber threat information. In addition, as threat information sharing platforms mature in their capabilities, this will aid in reducing network operators' response time.

⁵⁹ Public websites include sites like Twitter, Amazon AWS, Google Cloud, and Rapidshare.

The key for any successful coordinated network management against botnets is close, trusted collaboration and communications between stakeholders.

Industry Recommendations

This paper sets forth some of the problems presented by bots and botnets and the challenges and opportunities facing the owners and operators of broadband networks. The following section focuses on the preliminary recommendations that may be actionable by not only network service providers but the entire internet ecosystem to help reduce the threats from botnets using existing technology. The preliminary recommendations here are from the CSCC's perspective. There is a need to discuss best practices and capabilities for all segments of the ecosystem including software developers along with cloud, hosting, and application infrastructure providers.

Attack Mitigation

- **Encourage continued migration to all IPv6.**

The broad use of IPv6 will allow devices to have a unique address and can make it easier to track down the source of malicious traffic under certain circumstances.

- **Ensure that shared cyber threat information is actionable and tailored to meet the needs of recipients.**

Cyber threat information that is shared between internet stakeholders needs to be actionable by the recipients. Information sharing pool participants should tailor the information they share with their peers to be actionable.

- **Include pre-negotiated provisions for traffic filtering in transit and peering agreements.**

Network service operators of all sizes (ISPs, enterprises, governments, educational institutions, etc.) and end-users should ensure they have provisions in place with their

internet transit providers and peering networks to provide for upstream filtering and scrubbing of malicious traffic.

- **Streamline the law enforcement botnet takedown process.**

Law enforcement can play a key role in neutralizing botnets. Efforts are necessary to streamline the law enforcement process to increase the speed and efficacy of law enforcement botnet takedowns.

- **Encourage ICANN, registries, and registrars to adopt the fast flux mitigation techniques in SAC 025 SSAC Advisory on Fast Flux Hosting and DNS.**

The internet ecosystem should encourage ICANN, registries, and registrars to consider and adopt the fast flux mitigation techniques in the SSAC advisory.

- **Adapt and apply machine learning to the detection of botnets.**

The internet ecosystem should move away from manually reverse engineering botnet domain generation algorithms and begin applying machine learning to automate the real-time detection of botnets using fast flux, encryption, and other techniques to mask their infrastructure.

Endpoint Prevention

- **Ensure all end-points including IoT devices adhere to industry developed security standards.**

Multiple industry-led efforts are underway to develop security standards for IoT devices. IoT device manufactures and IoT service providers should work to ensure all IoT devices adhere to their respective industry security standards and best practices for IoT security.

- **Ensure end-points are running up-to-date software.**

As the saying goes “an ounce of prevention is worth a pound of cure.” This applies to consumer/customer end-points as well. Ensuring that all end-points (desktops, mobile, IoT, etc.) are running up-to-date software with the latest security patches and updates

will help tremendously in reducing the number of infected and compromised end-points on the internet.

- **IoT devices should use network isolation and/or network-based filtering techniques for any communications to cloud-based services.**

Network isolation and/or network based filtering are proven techniques for reducing the ability of a rogue internet end-point from doing harm.⁶⁰ IoT device manufacturers and IoT service providers should design their products and services to make use of these techniques.

Conclusion

Cybersecurity is a shared responsibility. Securing the internet from threats from botnets requires the collaboration and cooperation of all members of the internet ecosystem, both domestically and internationally. The preliminary recommendations in this paper represent just some of the many ways that botnet threats and their capacity for harm can be reduced through broad engagement by the stakeholders.

About the Authors

Matt Tooley is the Vice President of Broadband Technology at NCTA – The Internet and Television Association. He is a member of the Communications Sector Coordinating Council's Executive Committee. Tooley has over 30 years of experience in the broadband industry in developing and deploying broadband technology for internet service providers.

This paper includes key contributions from AT&T, CenturyLink and Cox Communications.

⁶⁰ BITAG, *Internet of Things (IoT) Security and Privacy Recommendations* (Nov. 2016) at Sec. 6 (discussing “A possible role for in-home network technology”), available at [http://bitag.org/documents/BITAG_Report_-_Internet_of_Things_\(IoT\)_Security_and_Privacy_Recommendations.pdf](http://bitag.org/documents/BITAG_Report_-_Internet_of_Things_(IoT)_Security_and_Privacy_Recommendations.pdf) (accessed June 20, 2017).

Appendix A - Cyber Threat Reports

Top 10 Worst Botnet Countries		
Rank	Country	Number of Bots
1	China	1,375,637
2	India	958,814
3	Russian Federation	569,463
4	Brazil	429,942
5	Vietnam	380,639
6	Iran, Islamic Republic Of	242,909
7	Argentina	177,701
8	Thailand	173,027
9	Mexico	145,516
10	C?*	141,684

Source: Spamhaus as of June 29, 2017. <https://www.spamhaus.org/statistics/botnet-cc/>

* Spamhaus reports the tenth country on this list as "C?."

Top 10 Botnet Traffic Attacking Countries		
Rank	Country	Percentage of Attack Traffic
1	China	50.8%
2	South Korea	10.8%
3	United States	7.2%
4	Egypt	3.2%
5	Hong Kong	3.2%
6	Vietnam	2.6%
7	Taiwan	2.4%
8	Thailand	1.6%
9	United Kingdom	1.5%
10	Turkey	1.4%

Source: Incapsula Global DDoS Threat Landscape Q1 2017. <https://www.incapsula.com/ddos-report/ddos-report-q1-2017.html>

Top Countries by % of Countries' IP Addresses Participating in DDoS Attacks, Q1 - Q4 2016 ⁶¹							
Q1 2016		Q2 2016		Q3 2016		Q4 2016	
Country	% of Countries IP Addresses	Country	% of Countries IP Addresses	Country	% of Countries IP Addresses	Country	% of Countries IP Addresses
	Source IPs		Source IPs		Source IPs		Source IPs
Turkey	0.282%	Vietnam	0.130%	U.K.	0.036%	Russia	0.078%
	43,400		20,244		44,460		33,211
Brazil	0.075%	China	0.093%	Brazil	0.025%	U.K.	0.059%
	36,472		306,627		81,276		72,949
China	0.035%	Taiwan	0.081%	China	0.025%	Germany	0.042%
	115,478		28,546		81,276		49,408
South Korea	0.028%	Canada	0.026%	France	0.025%	China	0.014%
	31,692		20,601		23,980		46,783
U.S.	0.005%	U.S.	0.006%	U.S.	0.004%	U.S.	0.012%
	72,598		95,004		59,350		180,652

Sources:

Akamai's State of the Internet Security Q4 2016 report. <https://www.akamai.com/us/en/multimedia/documents/state-of-the-internet/q4-2016-state-of-the-internet-security-report.pdf>

Wikipedia contributors, "List of countries by IPv4 address allocation," *Wikipedia, The Free Encyclopedia*, https://en.wikipedia.org/w/index.php?title=List_of_countries_by_IPv4_address_allocation&oldid=776891748 (accessed July 17, 2017).

⁶¹ The number of source IPs participating in DDoS attacks is from the Akamai State of Internet Security Report Q4 2016 report. The data has been normalized for the percent of a countries' assigned IPv4 addresses from IANA data at the time of the writing of this paper. The percentages may vary some from the time of the Akamai report.

Appendix B – Threats from Botnets

Click Fraud

Websites are often paid for by advertisers. Advertisers pay by the number of “clicks” or visits to the advertiser’s website. If a website or advertising broker is able to generate a perception that many people are visiting an ad, it compels the advertiser to pay for each of those visits. One way to generate lots of clicks is to command a botnet to generate those visits.

Email spam, phishing email, or malware email

Botnets are often used to originate unsolicited bulk email, which may also include distribution of malware of various types such as ransomware, links to phishing sites, and malware associated with bots. Botnets can also be used to send more mundane unsolicited sales propaganda.

Unauthorized Network Gateway

Bots within a protected network boundary such as an enterprise network can become unauthorized gateways into the protected boundary, and can be used to gain access to other resources (data or computers) within the protected boundary (aka lateral movement).

Data Theft

Bots can steal data from infected devices through means such as network monitoring, key logging, or scraping data from memory or disk. This is frequently accomplished because many bot members sit within private and enterprise networks next to assets containing the valuable data. A great amount of data theft today is accomplished with “Spear Phishing”⁶² attacks where valid looking emails are sent to a person at a company and that email is used to steal intellectual property or banking information, or to host malware. A typical attack may consist of the “bad guy” sending an email to an administrative assistant or other lower level employee that looks like it came from a senior executive, whereby the “executive” is asking for the email recipient to reset a password because an “invoice needs to be paid” today. The recipient will reset the

⁶²Federal Bureau of Investigation (FBI), *Spear Phishers* (Apr. 1, 2009), available at https://archives.fbi.gov/archives/news/stories/2009/april/spearphishing_040109 (accessed July 17, 2017).

password using obfuscated links containing malware in the email. This allows the infection to begin and the installation of APT (Advanced Persistent Threat) software conducts illegal activities.

Illicit Content Distribution

Bots are sometimes connected to peer-to-peer file sharing networks to help store and distribute illegal content.

Brute force password guessing

Botnets are used for brute force password guessing. One method uses high speed password guessing attempts using a random password algorithm, a password dictionary or a predefined password list. First, brute forcing can be used by an individual bot member as a recruitment method to infect other devices by scanning for any assets with a known open exposed port and then implementing one of the brute force methods explained to “guess” the password. Second, it can be used by a bot or botnet to brute force an intended targets login credentials to gain access to the privilege or data the credential provides.

Processing Theft (e.g., Bitcoin mining)

Due to the number of bot members typically seen in botnets, and the rising price of crypto currency (*e.g.* Bitcoin), botnets are very frequently seen being used to “mine” for coins. The process for mining Bitcoins requires the solving of very complex math equations which when solved, award the miner a set number of coins. In order to be successful, a miner needs a tremendous amount of computing power to solve these equations in the least amount of time. This is where a botnet can be extremely useful. By harnessing the computing power of a larger number of bots and “commanding” those bots to act as miners, the botnet owner can use the combined processing of many bots to make Bitcoin mining very lucrative.

Botnets have also been used to harness the computing power of the infected devices in order to perform Bitcoin mining or other activities for the benefit of the malicious actors running the botnet and not the legitimate owners of the computing resources.

Glossary

AIS – Automated Indicator Sharing, The Department of Homeland Security (DHS) operates a free service for the exchange of cyber threat indicators.

Bot – A program that is installed on a system in order to enable that system to automatically (or semi-automatically) perform a task or set of tasks typically under the command and control of a remote administrator (aka bot master or bot herder).

Botnet – A network of internet-connected end-user computing devices infected with bot malware, which are remotely controlled by third parties for nefarious purposes

Command & Control (C&C) – A remote computer used to coordinate the actions of bots.

CTI – Cyber Threat Indicator is the information that is necessary to describe or identify an attribute of a cybersecurity threat.

DDoS – Distributed Denial of Service attack is an attempt to make an online service unavailable by overwhelming it with traffic from multiple sources.

DNS – Domain Name System is the hierarchical decentralized naming system for resources connected to the internet.

DNS Water Torture – An attack type where many end-points send queries for a victim’s domain with a random string prepended to the domain that overwhelms the victim’s authoritative DNS server and making the victim’s domain inaccessible.

DOTS – DDoS Open Threat Signaling is a method by which a device or application participating in DDoS mitigation may signal information related to current threat handling to other devices or applications.

ICANN – Internet Corporation for Assigned Names and Numbers is the nonprofit organization responsible for coordinating the maintenance and procedures the internet’s namespace.

IRC - Internet Relay Chat is an internet protocol that facilitates communicating in text using a client/server architecture.

IoT - Internet of Things is the umbrella term to reference the technological development in which a greatly increasing number of devices are connected to one another and/or to the Internet.

IPv4 – Internet Protocol version 4 is the fourth version of the Internet Protocol (IP). IPv4 is one of the core protocols and still routes most Internet traffic today.

IPv6 – Internet Protocol version 6 is the sixth version of the Internet Protocol (IP). IPv6 is the most recent version and was developed to address the anticipated problem of IPv4 address exhaustion. IPv6 is intended to replace IPv4.

Kill Chain – Idea put forth by Lockheed Martin to describe the phases of a targeted cyber-attack: 1) reconnaissance, 2) weaponization, 3) delivery, 4) exploit, 5) installation, 6) command & control, and 7) actions.

NAT – Network Address Translation is a method for remapping one IP address space into another by modifying the address in the IP packet headers to allow multiple end-points to share one address while they transit a network router.

Network Service Provider – A network service provider or operator is any enterprise that is operating a network that has an assigned autonomous system number (ASN).

Peering – Peering is the voluntary interconnection of two separated networks for the purpose of exchanging traffic between users on each network.

Peer-to-Peer (P2P) – Traditionally botnets clients communicate to a C&C server for commands. P2P botnets operate without a C&C server where each bot is both a client and a server.

Software Defined Networking (SDN) – An approach to computer networking that allows for the programmatic control of network behavior using open interfaces and decoupling the packet forwarding plane from the control plane to allow for the use of standard servers and Ethernet switches to provide the routing function instead of specialized routers.

SSAC – The Security and Stability Advisory Committee advises the ICANN community and Board on matters relating to security and integrity of the internet’s naming and address allocation systems.

Tarpit – A tarpit is computer that purposely delays incoming connections. It is a defensive measure to make spamming and network scanning slower. It is analogous to a tar pit in which animals can get bogged down and slowly sink under the surface.

Transit – Internet transit is the service of allowing network traffic to “transit” a network to reach another network. Small network operators and enterprises buy Internet transit to gain access the Internet.