

CenturyLink 2019 Threat Report

(FjxxFjg|jyn+0bxy~{.dyx

Sbyybj @OynpxYbghfYnof

Dziopnupo Tpnvctuj

RGVmZW5kZXJzIG9mIGEgQ2x1YW4gSW50ZXJuZXQ

Rgo Hsno. Gbhe Ouhe.

27cf535ce08009de860b43589bd16ba3

Lqwhuqhw Erxqfhu

The promise of Connected Security



Security shouldn't have to be so hard.

Chris Betz, CSO

Let me put it another way: Security can be complex. The true art is making security easy to use.

As a Fortune 150 company and the second largest U.S. communications provider to global enterprise customers, we are responsible for securing our own operations through a suite of hybrid IT, cloud, networking and communications solutions — in addition to those of our customers. As CSO for this company, I can attest to the fact that the pressures security leaders face today are many.

On one hand, we have the explosion of network traffic spurred by video, 5G, IoT, connected devices and a mobile workforce; on the other, we have a justified and growing intolerance by users — both internal and external — for anything less than always-on, flawless performance. Couple this with the patchwork nature of many of today's security solutions, which businesses are often left to stitch together on their own; the gap between security and engineering teams that often reflects security as an afterthought; and the shortage of qualified security professionals — and the picture can seem bleak.

But security *can* be simple: We believe that the inherent value of a security solutions provider should first and foremost be effective simplicity.

At CenturyLink, our security builds on two fundamental directives: to leverage our expansive global threat visibility and to act against the threats we see. Our unique and deep network-based threat intelligence makes our approach possible — and it is the foundation of Connected Security, our vision for seamless integration between security and the network to transform the communications of tomorrow.

The more we can do as a global security services provider to identify or impact malicious traffic before it hits our customers' infrastructure, the better customers can focus and prioritize their resources elsewhere. This is the promise of Connected Security and the premise upon which we have transformed our network into a threat sensor and proactive defense platform.



Disrupting the security threats that we face today — and the threats we will face tomorrow — requires more than intelligence. It requires a collective commitment to share what we see and to act on what we know. We look forward to continuing to work together as we drive toward simplifying security.

Sincerely,

A handwritten signature in black ink that reads "Chris Betz". The signature is stylized and cursive.

Chris Betz, CenturyLink CSO

Contents

Part I: The role of deep network-based threat intelligence..... 5

 Black Lotus Labs: Defenders of a clean internet 5

 Threat research and operations at scale8

Part II: Drilling down into the latest risks and attacks..... 14

 Combating botnets: The army of network-based threats..... 14

 A deep dive into mass malware research from Black Lotus Labs 15

 Network-based behavior of botnets 17

 Black Lotus Labs analysis of select mass malware families..... 19

 Monitoring DNS: An underutilized security measure22

 DNS tunneling: Obfuscating malicious activity22

 DNS hijacking: Redirecting domain queries24

 DGAs: Hunting high-entropy domains 26

 Entropy and new techniques..... 26

 The good, the bad and the ugly: How to defend against DNS-based threats.....28

 Disrupting the disruptors: Safeguarding against DDoS attacks29

 Be bold — but beware34

Conclusion.....35

Part I: The role of deep network-based threat intelligence

Staying secure in a rapidly changing threat landscape means seeing both the forest and the trees. A thorough understanding of bad actor motivations, knowledge of the scope and scale of global threats, researcher intuition — all are critical for the identification of key patterns, and all are resources that our threat intelligence teams leverage to fully assess the true impact of threats on the network. With log volumes growing exponentially and the increasing reliance on internet connectivity and the cloud, organizations face mounting security events that overwhelm internal security resources.

The internet is the business-critical platform on which we all live and communicate today. Given the global nature of the CenturyLink backbone and deep network peering, our vantage point affords us visibility into a large percentage of the world's internet traffic. Whether IP traffic originates or terminates on our network, or traverses across our core backbone, we can leverage traffic flow data and our threat intelligence to gain deep understanding of what is happening on the internet. Very few providers have the breadth and depth of global threat visibility that CenturyLink possesses, and because of this, we recognize that we have a responsibility to help defend and protect the internet.

Black Lotus Labs: Defenders of a clean internet

It's a common saying in the industry that threat actors only need to be right once to be successful. But at CenturyLink, we know that threat actors also only need to be predictable to be caught. By modeling threat behaviors, understanding motivations, using attacker techniques as a kernel for research and analysis and ultimately implementing disruption efforts, we have built one of the world's most advanced threat research teams — Black Lotus Labs.



At CenturyLink, we are marshalling our forces into a defense platform that allows us to see more threats — so we can stop more. Black Lotus Labs is a direct reflection of our dedication and investment in focusing on deep network-based threat intelligence. The team hunts, identifies and observes bad actors attempting to leverage malicious code, then reviews their TTPs to identify the infrastructure and the C2s they are using.

**During 1H19, Black Lotus Labs tracked
18,000+ C2s daily**

However, data analytics alone cannot find a proverbial malicious needle in a haystack.

Since 2013, Black Lotus Labs has been baselining the behavior of the CenturyLink global backbone by ingesting and analyzing billions of data records daily and then using this baseline to detect potentially malicious anomalies. Each day, machine learning models developed by Black Lotus Labs ingest over 139 billion NetFlow sessions and approximately 771 million DNS queries. Across the first half of 2019 (1H19), Black Lotus Labs tracked 3.8 million unique threats per month, on average. We correlate these tracked threats against our NetFlow and DNS metadata to alert customers to a potential compromise.

The threat discovery and validation done by Black Lotus Labs drives the fidelity of our deep network-based intelligence. Threat validation is part of a framework Black Lotus Labs developed to help ensure the quality of threat intelligence. The research team performs validation to enrich our intelligence with confidence that a suspected threat type was scrutinized and matches an expected output.



Black Lotus Labs' systems, on average, monitored for **~1.2M** unique threats daily during the first half of 2019.

These threats represent **~15M** distinct malicious indicators tracked during the same timeframe.

Breakdown of validated C2s by family - 1H19.

1,438 Mirai	1,394 Emotet
-----------------------	------------------------

1,240 Gafgyt	39 Xor_DDoS
------------------------	-----------------------

9 Necurs

For example, if Black Lotus Labs identifies or ingests a suspected C2 domain, the team attempts a handshake in the malware's proprietary protocol to validate that it is, in fact, a malicious C2. This enables the team to hunt, identify and disrupt global threats. Working with intelligence partners around the world, Black Lotus Labs validated 4,120 new C2s during 1H19, which equates to roughly 686 C2s per month.

CenturyLink threat visibility:

~139B
NetFlow sessions
ingested and analyzed
per day

~771M
DNS queries
collected
per day

Of the new C2s identified during the first half of 2019, Black Lotus Labs independently discovered and confirmed 1,935 of them, including 654 from Gafgyt, 622 from Mirai and 659 from Emotet.

As a defender of the internet, Black Lotus Labs starts with the ability to detect and identify adversaries. But the team doesn't just passively observe malicious actors. Black Lotus Labs demonstrates a willingness to act by impacting the ability of verified malicious hosts to access the CenturyLink backbone and the internet.

Once the team gains high confidence that a host is acting as a malicious C2 and that removing it will make an impact, Black Lotus Labs works with the identified upstream service providers to disable the malicious infrastructure. During the first half of 2019, Black Lotus Labs notified other service providers of 468 verified C2s per month, on average, based on malicious activity witnessed. If providers do not take action, CenturyLink does so by removing the ability of approximately 63 C2s per month to access or send data across the CenturyLink global network.

Black Lotus Labs takes down **~63 C2s** per month from the CenturyLink network

Identifying, verifying and removing C2s from the CenturyLink backbone demonstrates that we are not passive observers with respect to malicious activity but are driving change on the internet, thereby helping protect our network and downstream customers.



Threat research and operations at scale

Organizations use a variety of methods to identify and track botnets, malicious actors and compromised hosts today. For example, honeypots are used to identify scanning behavior as well as to identify new malicious payloads being delivered through exploits. Internet-wide scanning can help identify C2s through headers, responses or SSL certificates. Malware analysis can enable security practitioners to understand the infrastructure with which malware samples may be communicating. Researchers use domain sinkholes to redirect a malicious domain name to their own controlled IP address, allowing them to identify all the bots that are querying that domain. Both open-source and third-party paid feeds are also used to track IoCs related to generic threat types or specific malware families.

Black Lotus Labs uses all these approaches and adds this information to a custom-built reputation system. The team ingests about 7.1 million unique entities per day, and at any given time considers 1.2 million of these to be active threats. The term “entity” is used to signify an IP address, domain or malware hash. Because CenturyLink is one of the largest internet backbone providers in the world, the team correlates the 1.2 million entities with NetFlow and DNS data. Black Lotus Labs then runs machine learning and heuristic-based algorithms on top of these correlations to find suspected malicious infrastructure related to these entities.

The metrics on the following pages represent the threats monitored by Black Lotus Labs over time and are broken down by threat type and suspected country of origin. The team determines country of origin by taking the IP address of each host and comparing it against a rich set of IP addresses to geographical mappings. The tracked daily averages for C2 hosts, suspected botnet hosts, hosts issuing attack commands, hosts distributing malware and hosts scanning the internet for vulnerabilities during 1H19 are shown in the following charts.

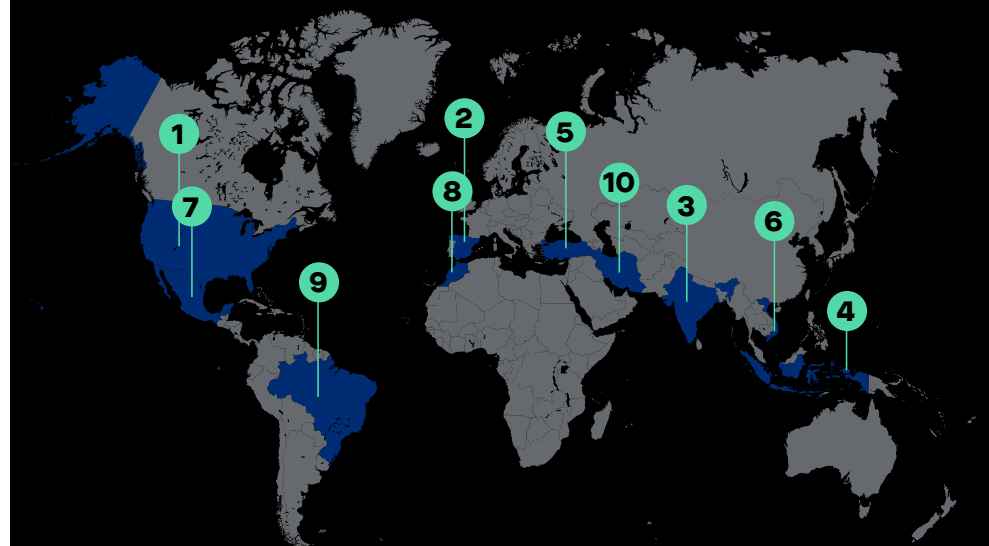
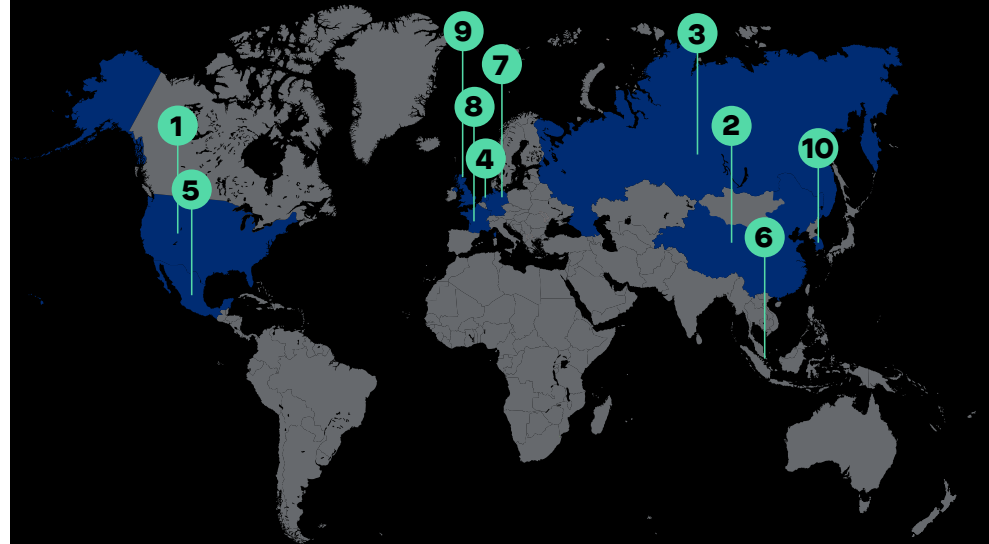
Global threats tracked by country, daily average

C2 hosts

1. United States	59.7 K
2. China	13.7 K
3. Russia	4.4 K
4. Netherlands	3.7 K
5. Mexico	3.2 K
6. Singapore	2.7 K
7. Germany	2.4 K
8. France	2.3 K
9. United Kingdom	2.0 K
10. Korea	1.4 K

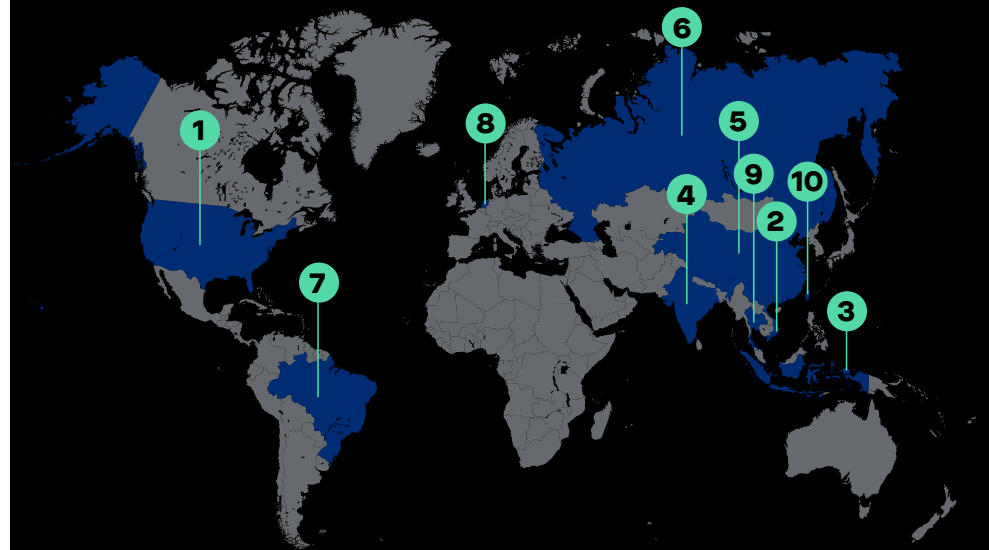
Suspected botnet hosts

1. United States	93 K
2. Spain	76 K
3. India	75 K
4. Indonesia	43 K
5. Turkey	32 K
6. Vietnam	32 K
7. Mexico	28 K
8. Morocco	26 K
9. Brazil	25 K
10. Iran	24 K



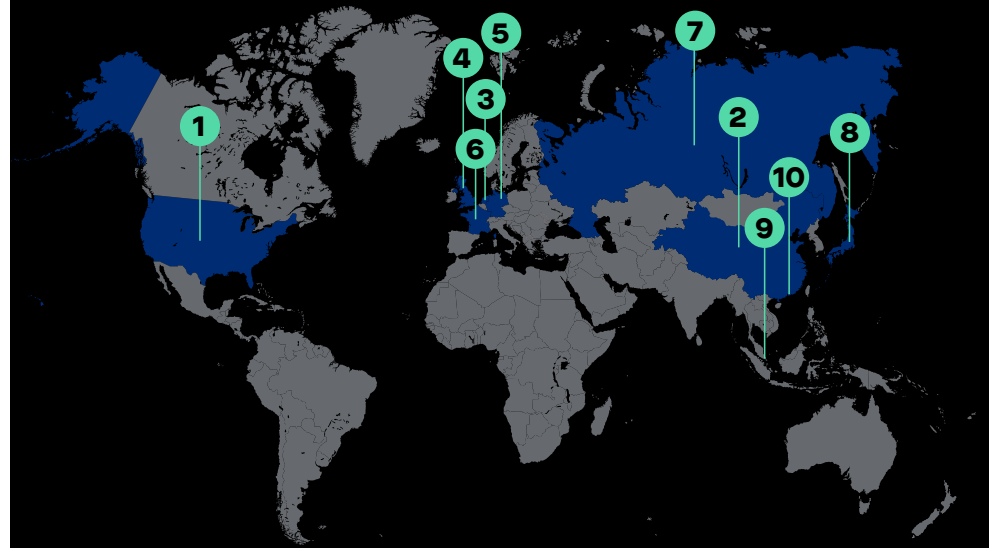
Hosts issuing attack commands

1. United States	45 K
2. Vietnam	27 K
3. Indonesia	17 K
4. India	14 K
5. China	13 K
6. Russia	13 K
7. Brazil	7 K
8. Netherlands	7 K
9. Thailand	7 K
10. Taiwan	6 K



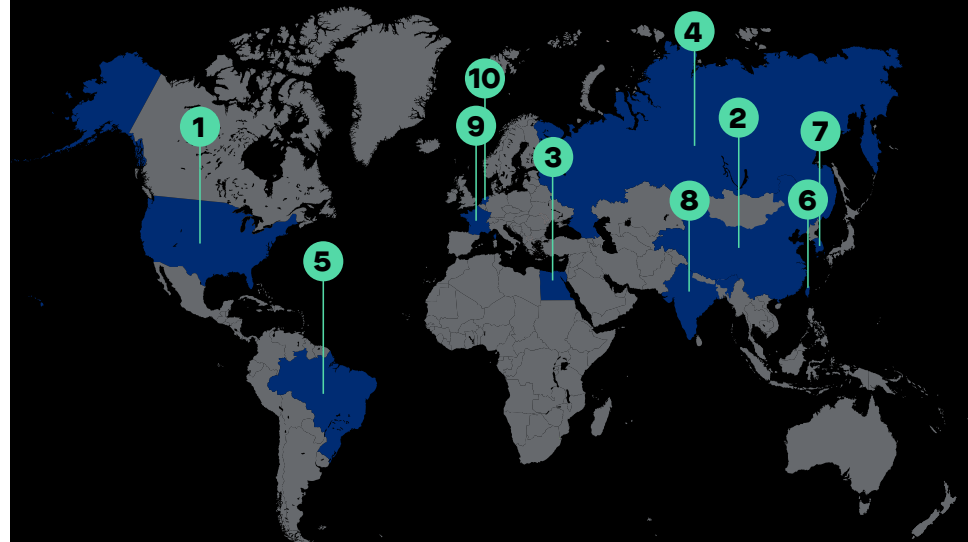
Hosts distributing malware

1. United States	177 K
2. China	54 K
3. Netherlands	6 K
4. United Kingdom	4 K
5. Germany	3 K
6. France	3 K
7. Russia	2 K
8. Japan	1 K
9. Singapore	892
10. Hong Kong	836



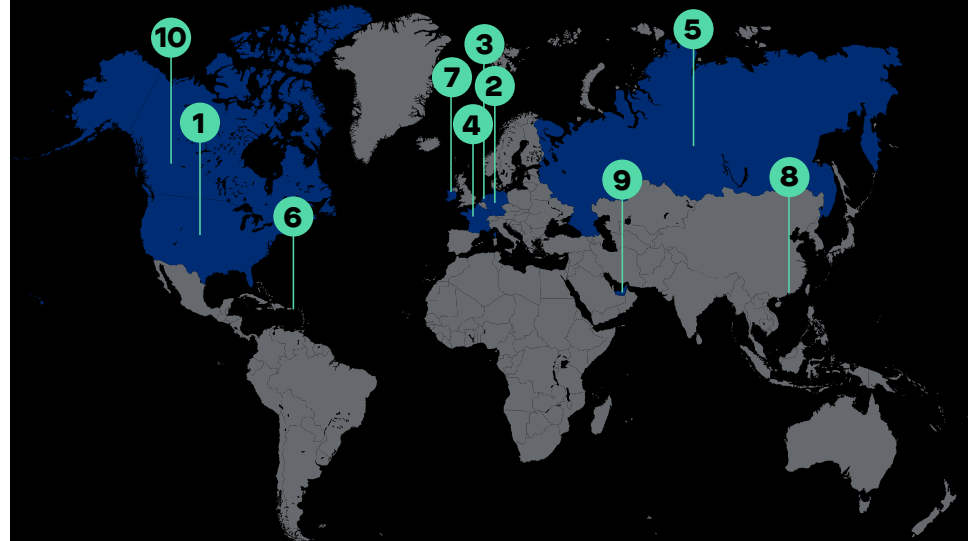
Hosts scanning for vulnerabilities

1. United States	176 K
2. China	87 K
3. Egypt	65 K
4. Russia	58 K
5. Brazil	53 K
6. Taiwan	35 K
7. Korea	27 K
8. India	24 K
9. France	21 K
10. Netherlands	19 K



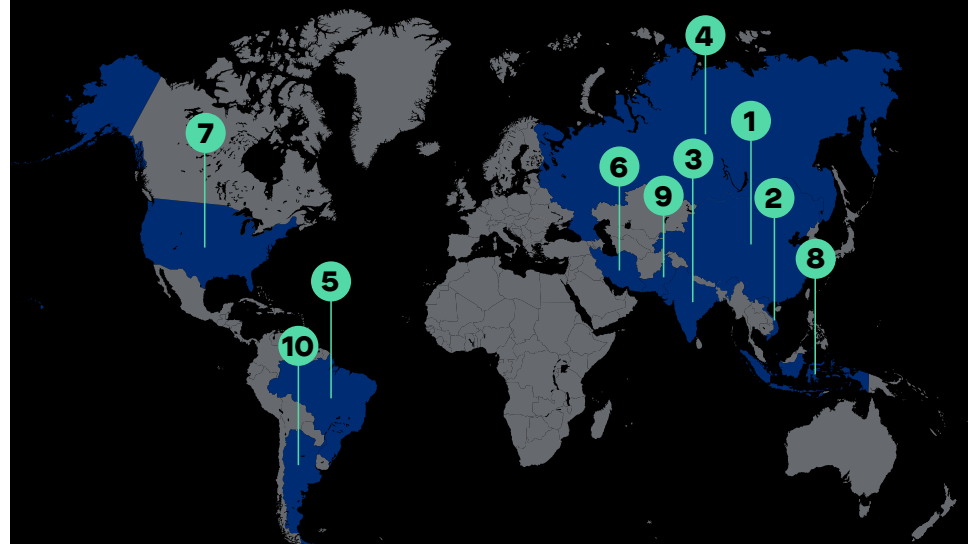
Hosts containing phishing sites

1. United States	202 K
2. Germany	29 K
3. Netherlands	21 K
4. France	16 K
5. Russia	15 K
6. Virgin Islands	10 K
7. Ireland	8 K
8. Hong Kong	7 K
9. United Arab Emirates	6 K
10. Canada	5 K



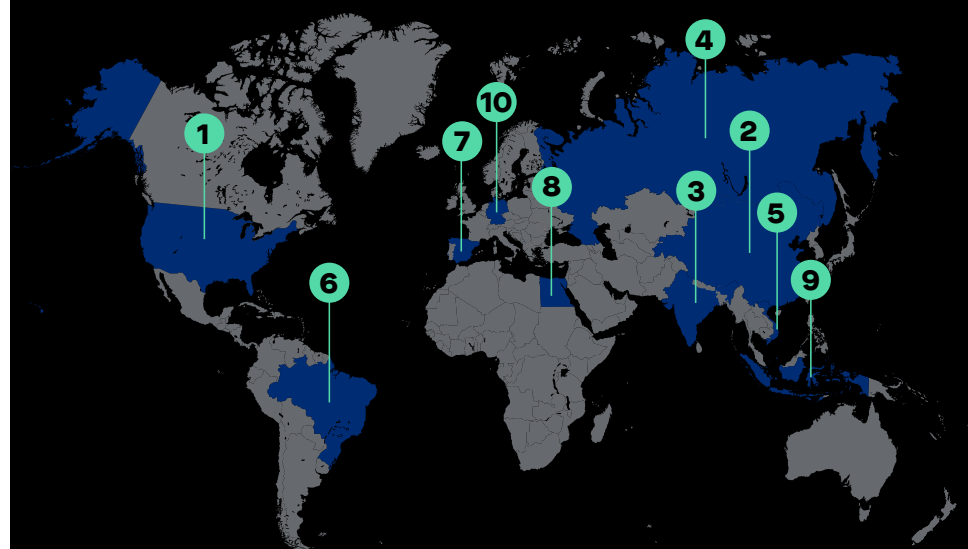
Hosts sending spam

1. China	103 K
2. Vietnam	61 K
3. India	34 K
4. Russia	30 K
5. Brazil	30 K
6. Iran	14 K
7. United States	14 K
8. Indonesia	13 K
9. Pakistan	13 K
10. Argentina	11 K



All threats: Cumulative threats tracked by country

1. United States	766 K
2. China	290 K
3. India	152 K
4. Russia	146 K
5. Vietnam	136 K
6. Brazil	117 K
7. Spain	94 K
8. Egypt	94 K
9. Indonesia	84 K
10. Germany	60 K



Black Lotus Labs also operates an extensive network of honeypots known as a honeynet. Each honeypot contains a variety of seemingly valuable, but innocuous resources for threat actors to attack and attempt to manipulate. Black Lotus Labs logs attacker TTPs for analysis and to produce actionable and trusted threat intelligence. The team then adds this intelligence to their reputation system.

A review of honeypot logs over 1H19 indicates that **~2.5M** unique hosts attempted to connect to the Black Lotus Labs honeypots, with **~33K** hosts making daily contact.

This equates to approximately **764** events per minute — or about **12** every second.



The volume of attempted interactions is compounded by both the ongoing prevalence of IoT botnets and the speed at which adversaries are developing new techniques to capitalize on existing vulnerabilities and to compromise new devices, as well as the infrastructure of other actors.

If you think of identifying malicious infrastructure as finding a needle in the haystack of the internet, Black Lotus Labs' machine learning-based algorithms reduce the haystack to a handful of hay. To find the needle, automated threat validation as described above helps to determine whether a host is indeed a threat. Validated threats are then fed back to the reputation system, allowing Black Lotus Labs to track mass malware families over time.



Part II: **Drilling down into the latest risks and attacks**

Combating botnets: The army of network-based threats

As companies focus on digital innovation, they are entering a new world of unprecedented threat and risk that continues to evolve. The lone-wolf troublemakers and attackers motivated by chatroom fame have been replaced by well-financed nation-states and criminal actors willing to play the long game and invest heavily in attack strategies and techniques. Threats continue to evolve and have become increasingly sophisticated — hypertargeted spear-phishing attacks, fileless malware, cryptojacking, chatbot attacks and IoT DDoS attacks are the current state. Cybercrime as a service isn't coming — it's already here.

Botnets exist because criminal actors must operate at scale to maximize profit

Two reasons for the continued success of botnets are the ease with which they compromise their targets and their ability to be operated remotely and covertly by cybercriminals. Many devices, from home security cameras to smart appliances, continue to be desirable targets because they are easy-access devices with limited security features — and consumers have been slow to realize and respond to their contribution to global cybercrime. Initially these connected devices were used for low-level attacks, but at CenturyLink today, we see a proliferation of connected devices being used for malicious infrastructure by more advanced actors.

Botnets exist because criminal actors, in order to maximize profit, need to operate infrastructure at scale. They are used today for every kind of malicious activity you can imagine — from sending massive amounts of spam and launching large-scale attacks that take down popular services to malicious hacking attempts. Some botnets are even sophisticated enough to recognize when they've breached a high-profile or otherwise valuable network and, in turn, sell access to other actors for their own malicious intent.



At CenturyLink, we recognize the importance of disrupting the work of known bad actors to help protect our customers' interests and preserve the security of the internet. Is it a bit of a cat-and-mouse game? Absolutely. However, these efforts are by no means fruitless. We believe that our actions help make it that much more difficult and costly for bad actors to operate, and we will continue to work to the best of our ability to prevent malicious actors from amassing infrastructure capable of disturbing the internet at large.

A deep dive into mass malware research from Black Lotus Labs

In the [CenturyLink 2018 Threat Report](#), we shared a deep dive into the evolution and presiding trends with respect to two widespread IoT DDoS botnets: Gafgyt and Mirai. At the time of the report, we noted that while Mirai garnered an inordinate amount of media attention, Gafgyt was significantly more prevalent based on our visibility.

This year, we revisited Gafgyt and Mirai, adding Xor_DDoS, which shares the same preference for Linux-based devices. Based on Black Lotus Labs' 1H19 data, we're seeing some notable differences since the last report. For example, the number of unique C2s for Gafgyt and Mirai for Q1 2019 alone appears to be on pace to far surpass those noted in the 2018 report. For all of 2017, Black Lotus Labs tracked 562 unique Gafgyt C2s and 339 Mirai unique C2s. In contrast, the current figures in the following chart represent just the first half of 2019.

Black Lotus Labs IoT malware family visibility — 1H19

	Gafgyt	Mirai	Xor_DDoS
Unique C2s tracked	676	989	20
Unique victims targeted	61,911	88,114	9,575
Average uptime (days)	12	14	32

Although there were fewer unique C2s reported for Xor_DDoS during 1H19, the average uptime is measuring nearly double that of Gafgyt or Mirai. There are a number of potential reasons why average uptime from Gafgyt and Mirai is decreasing:

- More researchers and threat teams are tracking their movement from upstream providers to new carriers.
- Black Lotus Labs automated the process of alerting upstream providers of C2s located on their networks.
- Threat researchers are getting better at identifying new Gafgyt and Mirai variants that were created to evade detection.
- Providers are getting better about proactively identifying and tracking C2s on their networks rather than reactively addressing the issues after resources have been consumed.

Network-based behavior of botnets

Over the last 12 months, Black Lotus Labs performed and published significant analysis on several botnets, which helps network defenders understand how mass malware campaigns operate. The following sections outline specific details of network behavior from select botnet families observed by Black Lotus Labs.

Mylobot

The Mylobot malware is particularly noisy from a network perspective. Every day, each bot attempts to resolve 1,404 hard-coded random-looking domain names. For each domain, the bots send out queries to resolve 43 subdomains, for a total of 60,372 DNS queries from each bot. If one of the queried domain names resolves to an IP address, the malware tries to connect to the IP address on a specific port that was hard-coded with that domain. When a bot receives a response from the C2, it XORs it with the byte 0xDE. The resulting string contains up to two URLs. Each URL contains an IP address and a file ending in .gif. The malware connects to this URL and executes the downloaded file the team identified as Khalesi — an information stealing malware. The bot then executes the downloaded file. This flexibility with the downloader allows the malware authors to modify the second stage whenever they choose.

In late 2018, Black Lotus Labs saw 7,764 distinct IP addresses querying these 1,404 hard-coded domains, and a total of 13,402,925 DNS queries in CenturyLink's passive DNS data. The two files Black Lotus Labs saw in the C2 response were PE32 executables.

TheMoon

TheMoon is a modular botnet that targets various vulnerabilities in IoT devices. C2 communication from the initial payload is done on three different TCP ports, typically in the 4000 to 6000 range. From a network perspective, TCP communication to a single IP address on multiple nonstandard ports in this range may be indicative of C2 communication. TheMoon uses each port for different purposes: one for registration, one for command and control, and one for downloading payloads. Payloads are secondary binaries that have their own C2 communication, sometimes to the same IP address as the main module in a similar range of ports. This malware family is known to turn devices into proxies for use by other actors. CenturyLink recommends checking internet-connected devices to ensure that no unknown ports or services are open. Proxy ports for TheMoon are above 10,000 and appear ephemeral in nature. At its peak during mid-2018, TheMoon averaged over 10,000 proxies per day.

Necurs

Necurs bots communicate with their C2s using HTTP POST requests on port 80. The URL path has always been a PHP file in our observations, and the host header always uses an IP address rather than a domain. Depending on the modules installed by the bot, it may communicate with C2s on other ports as well, such as port 5222. Bots will send a POST request to the C2 with an encrypted payload of around 250 bytes, which includes information about the bot host, such as its Windows version and its bot ID. When attempting to find a new C2 via its peer-to-peer network, it will send 29-byte encrypted payloads to all known peers, primarily over UDP, but sometimes using TCP, using the same source port for each peer. When trying to find a new C2 via its DGA, it will query up to 2,048 DGA domains, sending each a 50-byte payload via HTTP POST requests. The IP address it receives when resolving its DGA domains will be de-obfuscated by the bot, and the bot will send the POST request to the de-obfuscated IP address rather than the one in the DNS response. Despite having been resolved from a domain, the POST request will still use the de-obfuscated IP address in the host header, rather than the DGA domain.

Keeping tabs on Necurs



A global logistics company implemented CenturyLink’s security services to review its network security logs, which revealed a single user with logins into 65 separate accounts. Using CenturyLink threat intelligence, the user’s behavior was identified as that of an infected Necurs host. CenturyLink immediately notified the customer, and the issue was remediated less than one hour from initial contact. [Watch the video.](#)

Black Lotus Labs analysis of select mass malware families

<u>Emotet</u>	<u>Mirai/Satori</u>	<u>Mylobot</u>	<u>Necurs</u>	<u>TheMoon</u>
What it is				
After emerging in 2014 as a prominent banking trojan, Emotet morphed into a modular spam and malware-as-a-service botnet with global distribution. It primarily targets Windows machines.	Mirai debuted in 2016 as malware developed by feuding video game server operators. Because it provided an easy-to-use, scalable framework for IoT attacks, it quickly became popular as a generic DDoS attack platform. Since that time, Mirai malware has evolved into several variants that are capable of amassing internet-connected devices to launch similar attacks. Satori is one variant of the Mirai bot.	Mylobot emerged in June 2018 as sophisticated malware that uses interesting techniques in an attempt to avoid detection. It functions as a downloader and is able to deliver varying types of other malware. In 2019, Black Lotus Labs observed it downloading the information-stealing malware, Khalesi.	First discovered in 2012, Necurs is one of the most prolific spam and malware distribution botnets in history, having infected many millions of computers. Necurs has evolved from a spam botnet delivering banking trojans and ransomware to a multifaceted tool capable of proxying traffic, enabling crypto-mining and launching DDoS attacks.	Identified in 2014, TheMoon is a modular IoT botnet that targets vulnerabilities in routers within broadband networks, such as those used by consumers and small businesses. Bad actors use TheMoon for credential-stuffing attacks, video advertisement fraud, internet traffic obfuscation and more.
How it infects				
Emotet typically spreads via malicious links or attachments in phishing emails sent from infected devices. As of May 2019, Proofpoint reported that Emotet was responsible for more than half of all spam containing malicious payloads at the time.	Mirai scans the internet for unprotected IoT devices with default usernames, passwords or exploitable vulnerabilities. After it connects to a compromised host, Mirai installs malware, thereby adding the vulnerable IoT devices to the overall botnet. Satori is an example of a variant targeting Android devices. In July 2018, CenturyLink found Satori bots attempting to infect such devices.	The initial infection vector eludes most researchers. However, after the initial compromise, Mylobot can instruct the compromised host to download other types of malware.	Links to Necurs malware are distributed through phishing emails, such as Russian dating scams and pump-and-dump stock scams. Users click links that initiate the download of the malware. This simplistic malware delivery method has proven successful. Necurs successfully infects computers running pirated or unpatched versions of Windows that typically don't have antivirus protection. As such, it is seen primarily in developing nations.	Operators of TheMoon scan for hosts by looking for vulnerable services running on unsecured IoT devices. Once a service is found, TheMoon uses an exploit to drop a shell script that, when executed, downloads the initial stage payload from domstates[.]su. TheMoon currently exploits broadband modems or routers developed by Linksys, Asus, MikroTik, D-Link and, most recently, GPON.

<u>Emotet</u>	<u>Mirai/Satori</u>	<u>Mylobot</u>	<u>Necurs</u>	<u>TheMoon</u>
Behavior				
<p>Although Emotet phishing campaigns have only been observed delivering Emotet itself, once a computer is infected with Emotet, it can then be used to drop additional malware families. Emotet's C2 infrastructure is complex. For example, it utilizes several tiers of hosts to control its infrastructure. Additionally, during certain infections, Emotet deploys a Universal Plug and Play (UPnP) module, which allows an infected device to act as a C2.</p>	<p>The Mirai C2s issue DDoS attack commands to each of their bots. These commands identify a target, which port to attack, which protocol to use and the duration of the attack.</p> <p>In January 2018, the Satori variant began targeting Claymore cryptocurrency miners, hijacking miners' wallet addresses to steal mined currency.</p>	<p>Mylobot contains sophisticated techniques designed to avoid virtual machines and sandboxing environments in an attempt to avoid detection. For example, the malware can instruct the host to lie dormant for 14 days before attempting to contact a C2.</p>	<p>Necurs uses a DGA to allow for a dynamic operation and avoid full takedown. When a compromised host can't reach its normal infrastructure, a predefined algorithm in the code generates a list of domain names that may be answering as a new C2. The host then attempts to connect with that C2. On occasion, the C2 network does seem to go offline for a period to avoid detection. During this time, the compromised hosts access the DGA in the malware and continue to look for a C2 to connect to.</p>	<p>TheMoon distributes malicious modules of differing functionality after initial infection. One of the modules attempts to create a proxy network for the operator's customers to use. Some of those customers may use the proxy as a service to send attack traffic, making it harder to track down the true source of the attack.</p>
Characteristics				
<p>Emotet's actors maintain over a hundred different C2s at any given time, frequently updating which C2s are active throughout the day. The C2 tiering structure noted above makes Emotet's infrastructure notably resistant to individual C2 disruptions.</p>	<p>Mirai botnets can each have only one C2, so when that is taken down, the botnet becomes orphaned. Orphaned bots often are reidentified and resubscribed to a new C2, if the initial compromise is unresolved.</p>	<p>Black Lotus Labs observed approximately 18,000 unique IP addresses communicating with Mylobot C2s. The bot performs DNS lookups that appear to be DGAs but that are actually predefined domains in the malware.</p>	<p>The bot uses a distributed peer-to-peer network to connect infected machines.</p>	<p>TheMoon has several hard-coded IPs it uses for C2 communication within its main binary. Secondary payloads may have separate C2 infrastructure.</p>

<u>Emotet</u>	<u>Mirai/Satori</u>	<u>Mylobot</u>	<u>Necurs</u>	<u>TheMoon</u>
Payloads				
Over the past several years, Emotet has mainly been used as a malware distribution service for other families such as Trickbot, IcedID, QakBot, Gookit and Dridex, and it has become central to the cybercrime ecosystem.	Mirai does not drop a secondary payload.	Mylobot is able to deliver several secondary payloads. One example is Khalesi, an information-stealing malware. Khalesi is a keystroke logger that steals financial and other data.	Prior to payload delivery, Necurs is able to evaluate hosts to determine whether they qualify as potential botnet members. For this reason, it will not deliver secondary payloads in a lab environment or attack older-model computers that lack computing power.	Similar to Mylobot, TheMoon can distribute additional secondary payloads of varying functionality. Most secondary payloads are associated with turning the bot into a proxy offered as a service.

How to defend and protect				
Since Emotet is distributed via phishing emails, users should be trained not to open unknown email links or attachments. Emotet's binary is typically distributed via Microsoft Word macros. Training users to not enable macros or disabling macros altogether is also recommended. Since Emotet's Word macros typically run a script via Windows PowerShell, we recommend that businesses use endpoint monitoring to detect the execution of Windows PowerShell from Word. Emotet's C2 servers are freely available from several sources online, and they can be used for detection or blocking purposes at the network perimeter.	Companies should be aware of all devices accessing the internet through their network and baseline that traffic in an effort to flag anomalies. In addition, companies should adjust firewall configurations to minimize vectors that can be used to compromise internet-connected devices.	Organizations that monitor DNS can detect Mylobot as it attempts to contact a C2. As of November 2018, Black Lotus Labs identified 60,000 queries sent to hard-coded domains. Review logs to identify any spikes in DNS queries. Use DNS-based reporting and filtering to identify and block threats and receive analytics on attempted attacks.	Security teams should patch software proactively to avoid gaps that Necurs can exploit. They should also train users not to open unknown emails or click on unknown links. Security researchers can use methods like sinkholing of identified DGA domains along with analysis of DNS and network traffic to enumerate Necurs bots and C2 infrastructure.	IoT device manufacturers and broadband network providers should limit services open to the internet and continually provide patches for security vulnerabilities. Default usernames and passwords should be immediately changed when the device is first turned up.

Monitoring DNS: An underutilized security measure

Have you checked your DNS lately? Over the nearly 40 years DNS has been in use, everyone has become accustomed to leveraging DNS to navigate from domain name to IP address.

Yet as threat actors refine and evolve their TTPs, monitoring DNS traffic for malicious activity can be an effective security mechanism for today's defenders. By paying attention to DNS servers, organizations can leverage log data to track threats as adversaries ramp up DNS-based attacks to support a variety of malicious campaigns. Later in this report, you'll find an example of this approach by one of Black Lotus Labs' research partners, Cisco Talos.

Because CenturyLink operates one of the largest DNS resolution services on the internet, Black Lotus Labs possesses unique visibility to potentially identify and enumerate malicious infrastructure changes soon after they propagate. DNS monitoring allows the team to track advanced threats and describe actor behavior to correlate attack campaign timelines.

For this report, Black Lotus Labs focused on how DNS data can reveal important information about the ways in which actors operate DNS-based attacks. The next section details some common ways in which attackers take advantage of DNS to achieve their goals, including DNS tunneling for data exfiltration, DNS hijacking and the creation of DGAs.

DNS tunneling: Obfuscating malicious activity

Over time, defenders have gotten better at identifying and blocking malicious network traffic; however, certain types of traffic still go unnoticed. As a result, malicious actors have found ways to use well-known, rarely monitored protocols — like DNS — to their advantage. Though many organizations don't think of DNS tunneling as an attack vector, it can be used to encode data in subdomains of a DNS query or response.

Many organizations don't restrict what queries can be run within their environments. In many cases, organizations are content filtering and man-in-the-middle proxying HTTP traffic, but letting DNS traffic through. This can make for a dangerous scenario, as it allows anyone to send arbitrary traffic outside of the network. Although some may question the risk of the situation, it is the querying host that queries the authoritative DNS server, allowing for a direct communication path when the query is performed correctly.

DNS tunneling could be indicative of data exfiltration by sophisticated actors, or it could be used to subvert security controls by the placement of data in the query. This data could be host-level data of compromised machines on a company's network or something much more valuable to the attacker. Defenders should more closely monitor their DNS infrastructure and look for DNS-based anomalies that could indicate malicious activity.

Black Lotus Labs is conducting research into DNS tunneling and has built a detection model to identify possible encoded data in DNS queries. The model looks at the number of queries to a second-level domain and examines the subdomains being queried. Below are two examples of queries that were flagged as possible tunneling in TXT and NULL record types.

APT32, also known as OceanLotus, used C2 communication via DNS in the SOUNDBITE malware. The model flagged similar queries in passive DNS (PDNS) data:

- 72jjbQAAAAAAAAAAAAAAAAAAAAAAJS4.z.nsquery[.]net
- JCx0-wAAAAAAAAAAAAAAAAAAAAAAhz.z.teriava[.]com
- KBb5wQAAAAAAAAAAAAAAAAAAAAAAPML.z.tulationeva[.]com

The DNS queries consisted of NULL and TXT record types and included base-64 encoded data in the subdomains. This detection model also flagged four other domains that appeared to be encoding data in a similar format:

- SG1EFwAAAAAAAAAAAAAAAAAAAAAGe5.z.jessicajoshua[.]com
- SG1EFwAAAAAAAAAAAAAAAAAAAAAKrk.z.thomaswaechter[.]com
- G9PNggAAAAAAAAAAAAAAAAAAAAAMxF.z.phillippcliche[.]com
- G9PNggAAAAAAAAAAAAAAAAAAAAAPWc.z.poppyranken[.]com

Black Lotus Labs also observed DNS queries to a domain employed by Xshell Ghost malware, which uses a backdoor that was discovered in Xshell and leverages DNS TXT messages to exfiltrate data. The malware uses a DGA that generates a new domain each month for the DNS communication. The team saw queries such as the following to the domain luvifolufwbsb[.]com, which is the domain employed in April 2019. (See the [netlab360](#) and [arXiv](#) reports, which provided the initial research seed.)

- qajajlyoogrmkclivhqbtgpbpmwlvvcjikdtkkgwfvdkbkoqfukognnvfuh.xcohnqqlmoxgqftfrcsbtpq.luvifolufwbsb[.]com

Even though these techniques were discovered several years ago, this research shows that the threat was not eliminated — and in some cases, we've even seen an increase in this activity, including activity by new, previously unreported domains.

As an example, over a multiweek period, Black Lotus Labs' algorithms found an average of 250 domains per day actively being abused for the tunneling of data. This represents over 70,000 lookups to each domain, demonstrating the volume of data that can be visible when investigating these threats.

Black Lotus Labs identified ~250 domains per day being abused for data tunneling, representing 70,000+ lookups to each domain

DNS hijacking: Redirecting domain queries

DNS hijacking, the practice of subverting DNS resolution by manipulation, can take many forms — from changing the DNS settings on devices to diverting every packet of data that leaves a company's network by redirecting the destination of the domain. DNS hijacking also recently expanded to include DNS record manipulation. DNS registrars can be targets of password brute-forcing, but the more popular instances of DNS hijacking lately have been a result of DNS administrators being phished and their legitimate credentials being used to perform the attack.

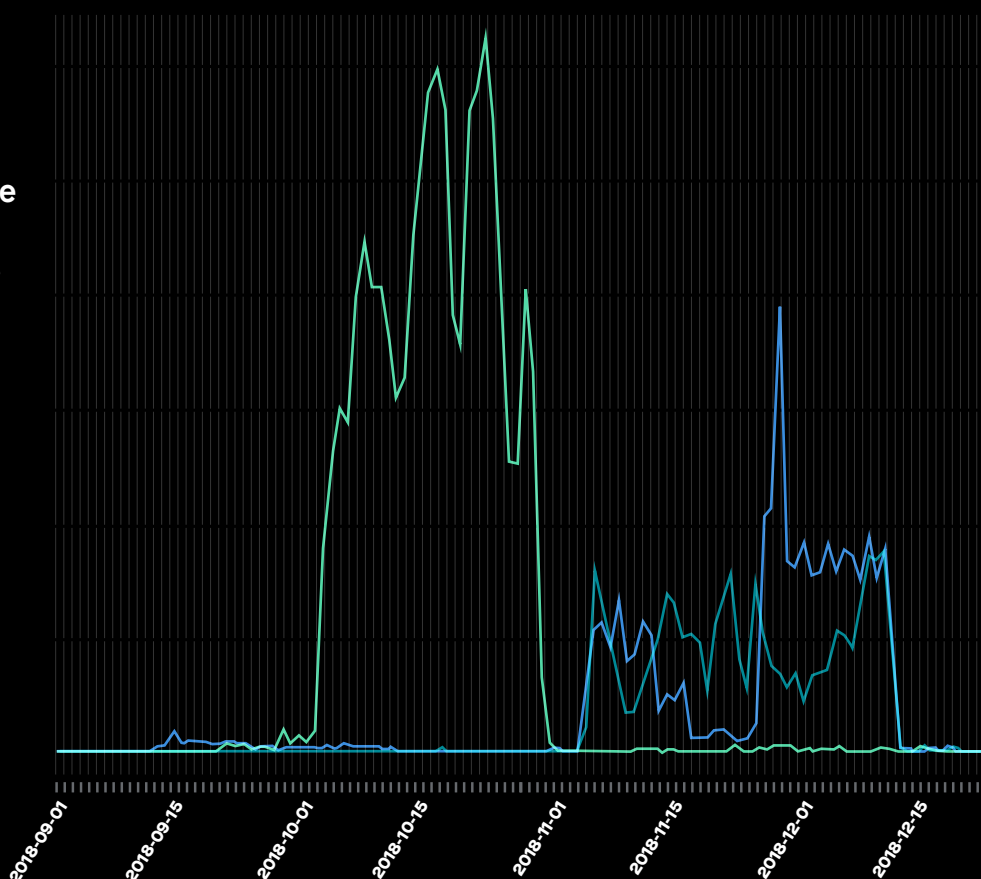
During the last quarter of 2018, one of Black Lotus Labs' closest threat intelligence partners, Cisco Talos, reported the emergence of an attack campaign with many pivots, which it called [DNSSpionage](#). Attackers phished usernames and passwords to compromise the DNS integrity of targets in one pivot, aiming to obtain registrar credentials. Through these credentials, the attackers installed new host names and had new TLS certificates issued to become valid users of the domain and impersonate the site. Even though the attackers pointed stolen domains to their malicious hosts for only a few minutes, they were able to obtain vast amounts of user and VPN credentials. This success affords the attackers time to develop custom malware and imposter websites, and to enumerate soft targets such as job postings in order to launch future targeted attacks under the radar. Although the DNSSpionage campaign went unnoticed for some time, when these brazen threat actors were caught in the act, they continued to launch DNS hijacking attacks, as also reported by [CrowdStrike](#) and [FireEye](#) earlier this year.

Black Lotus Labs continuously monitors for and tracks new threats to protect CenturyLink customers. In order to understand the impact of the events described in the Talos report more broadly, Black Lotus Labs performed an independent post-event analysis on three IPs involved in the DNSpionage attack as illustrated in the graph below. On September 15, 2018, there was a small spike in 185.20.187.8, which directly corresponds to the DNS hijack identified by Talos. In October, the team saw a large increase in traffic to 185.161.211.72. CenturyLink's PDNS data shows that this IP was associated with the C2 domain from October 8-30, 2018, directly aligning with this increase in NetFlow activity. Throughout November, the team saw a shift in traffic to 185.20.184.138 and 185.20.187.8, which also corresponds to the C2 domain resolutions from November 5-27, 2018.

Independent NetFlow analysis by Black Lotus Labs of the DNSpionage campaign identified by Cisco Talos

**Talos DNSpionage
traffic review by
Black Lotus Labs**

■ 185.20.184.138
■ 185.161.211.72
■ 185.20.187.8



Black Lotus Labs continues to monitor for this threat to protect the CenturyLink network and those of its customers.

DGAs: Hunting high-entropy domains

The use of DGAs in malware itself abuses the domain registration process and how DNS resolves domains. Some malware families use algorithmically generated domain names for locating their C2 infrastructure. Though the security research community is getting better at mitigating control infrastructure by using hard-coded domains or IP addresses, DGAs allow the malware to generate a new domain so the actor can change infrastructure, which can be difficult for security researchers to detect. Black Lotus Labs uses machine learning models to predict and anticipate these domains, which appear to be randomly generated. Based on this research, the team can understand the cadence generally followed by threat actors when they evolve malware and how they shift tactics and techniques to stay one step ahead.

Malware authors have used algorithmically generated domains for several years, allowing them to update the infrastructure they use without having to update the malware itself. The simplest form of these algorithms starts with a value (which may be hard-coded or derived from another source such as date/time) to seed a pseudo-random algorithm that chooses letters, numbers or a combination of the two to make up the domain name. Examples of DGA domains based on different seeds from the [Necurs malware family](#) are as follows:

```
DGA05: smjteudue[.]pw  
DGA09: vyosbpxjoffckyntpgk[.]pw  
DGA13: vdnpaavoyjoyrfmlejpe[.]pw  
DGA15: tquxdrapf[.]pw
```

Flagging these domains is fairly straightforward because they do not conform to predictable linguistic models, resulting in higher entropy and a greater likelihood of standing out to analysis techniques. The level of entropy with respect to domains corresponds to the randomness of the domain entry — the higher the level of randomness, the higher the level of entropy. Black Lotus Labs uses the level of entropy combined with several other features to flag random DGA domains.

Entropy and new techniques

Because entropy is such an effective method of detecting these randomly generated domain names, actors have moved on to new techniques. For example, several families of malware now use algorithms that choose words at random from a fixed word list or dictionary to make up the domain name, thus lowering their entropy. A simple DGA domain such as

QmxhY2tMb3R1c0xhYnMuaW8[.]io shows a high level of entropy and is seemingly random — however, in this case, the malware possesses the key that transforms this high-entropy domain into one that is accessible.

Detecting dictionary-based DGAs is a more difficult problem because the resulting domains have lower entropy and a character distribution similar to that of English words. Black Lotus Labs developed models to detect these domains by extracting and analyzing the association between the words that make up the domain. Examples of dictionary-based DGAs that were flagged by this model include **uponthank[.]net** and **rathersafety[.]net**. Both DGAs have been reported by another third-party source as related to the malware family Suppobox.

Random DGAs won't be going away completely, partly because they can be programmed in a couple of lines of code, whereas malware that uses a dictionary-based algorithm needs to either have the dictionary hard-coded or easily accessible. The Black Lotus Labs models use natural-language processing and other machine learning techniques to flag both random and dictionary-based DGA domains. The team runs the models against a dataset made up of daily in-house PDNS data, which includes, in part, mappings between domains and the IPs they resolve to, which Black Lotus Labs frequently leverages to track malicious infrastructure. The resulting output of this analysis process includes flagging roughly 130,000 domains per day that are potentially random DGA or dictionary-based DGA domains.

Black Lotus Labs flags **~130,000** potential random DGA or dictionary-based DGA domains per day

During 1H19, Black Lotus Labs observed 1,549 malware families utilizing 6,883,871 distinct domains. Of those malware families and domains, 52 families are using DGAs, which accounts for 6,076,207 of all domains observed. The number of domains is relatively high due to the proliferation of malware that utilizes a DGA to produce C2 domains that constantly shift.

The good, the bad and the ugly: How to defend against DNS- based threats

DNS tunneling and data exfiltration

- Review DNS logs inside your infrastructure and monitor for anomalies, large volumes of changes, and previously unseen domains.
- Monitor the network for large volumes of traffic going to single DNS servers. Monitor endpoints to identify processes that do not normally communicate over the network.
- Use an upstream provider that blocks malicious domain names.
- Leverage a network-based intrusion detection system/ intrusion prevention system (IDS/IPS) with exfiltration-specific signatures.
- Look for statistical anomalies on domain lookups, which can make these attacks stand out in DNS log sets. DNS tunneling tends to be very loud.

DNS hijacking

- Ensure that DNS administrators, or anyone with access to your organization's DNS registrar, use multifactor authentication on the registrar. Multifactor authentication helps mitigate successful phishing attempts.
- Constantly audit public DNS records to verify that they are resolving as intended.
- Search for any encryption certificates related to suspect domains and revoke any fraudulently requested certificates.
- Use DNSSEC to mitigate unauthorized changes to a domain, such as DNS cache poisoning on a local network.



It's important to keep in mind that mitigation techniques are contextual to an organization's environment, and defending against DNS-based threats is not limited to these recommendations.

But as DNS-based campaigns evolve, today's defenders can gain valuable insights and bolster defenses by including DNS monitoring as part of their security measures.

Disrupting the disruptors: Safeguarding against DDoS attacks

DDoS attacks have been around for decades and continue to be a popular weapon of choice for cybercriminals to target and overwhelm everything from a server to an entire network. Cybercriminals have targeted all industry verticals and government entities, causing service delays or taking business operations entirely offline. Many attackers launch DDoS attacks to test the response capabilities of intended targets.

CenturyLink global SOC's mitigated 14,000+ customer DDoS attacks during 1H19

Though the industry at large has seen and will likely continue to see an ongoing progression in attack sizes, CenturyLink also observed an increase in short-duration, lower-bandwidth attacks often lasting 30 seconds to one minute. Actors are motivated by the fact that some providers' DDoS mitigation solutions struggle to mitigate these bursting attacks. Defenders, depending on their customer environment and applications, may need to evaluate their risk tolerance for service-disrupting, short-duration attacks and consider always-on, as well as automated mitigation options with any on-demand mitigation solution.

During 1H19, CenturyLink global SOC's mitigated more than 14,000 customer DDoS attacks (~120 attacks per day) — including an increase across multivector and mixed application layer attacks. Additionally, as malicious actors continuously shifted attack types, CenturyLink regularly varied countermeasures to effectively mitigate the attacks. Tuesday, February 5, marked the highest attack day by volume during the first half of 2019; many of those attacks involved extortion and appear to have been coordinated.



The CenturyLink global SOC's also relied on near real-time intelligence, enrichment of data and visibility from Black Lotus Labs, which enabled more informed decisions on the type of countermeasures to deploy. We could apply these directly as controls into the CenturyLink network by dropping bad traffic at the peering point or network edge using BGP FlowSpec or by engaging in a C2 takedown. When undertaking C2 takedowns, our priority is to do so without inflicting collateral damage. We deployed such controls in conjunction with traditional DDoS scrubbing infrastructure, signatures and heuristics-based filters for enhanced mitigation.

Because of CenturyLink's global IP backbone, Black Lotus Labs has the advanced capability to take down C2s that produce DDoS attacks. The ability to remove the source of the C2 infrastructure is critical in combating IoT DDoS attacks. With the increasing number of internet-connected devices, combined with the growing internet throughput available in consumer homes (consider the 1Gb speeds and 5G wireless advances coming soon), malicious actors have a powerful incentive to gather an arsenal of IoT devices under the same C2 infrastructure. Whether the C2 is using bots to launch a DDoS attack, scanning the internet or enterprise environment, or spreading laterally from connected devices to other assets in the home or enterprise, IoT devices are attractive endpoints with which threat actors can strengthen their botnets.

Game over, DDoS

DDoS operators use botnets, like Mirai, to launch attacks. These botnets can also be used to provide cover for other malicious activity. For a gaming company that was about to launch a new game, CenturyLink proactively blocked a botnet's C2 infrastructure after receiving chatter of a pending attack. That move enabled the company to launch its game successfully without worrying about its gaming platform availability being disrupted by a DDoS attack.

DDoS attack size, duration and type – 1H19



of DDoS attacks mitigated

Daily =

~120

1H19 =

~14K



Median attack size =

4 Gbps



Average size of peak daily attack =

25 Gbps



Largest attack seen =

430 Gbps



Median attack duration =

44 Minutes



Average attack duration =

1 Hour, 51 Minutes



Longest DDoS attack =

9 Days, 4 Hours, 40 Minutes

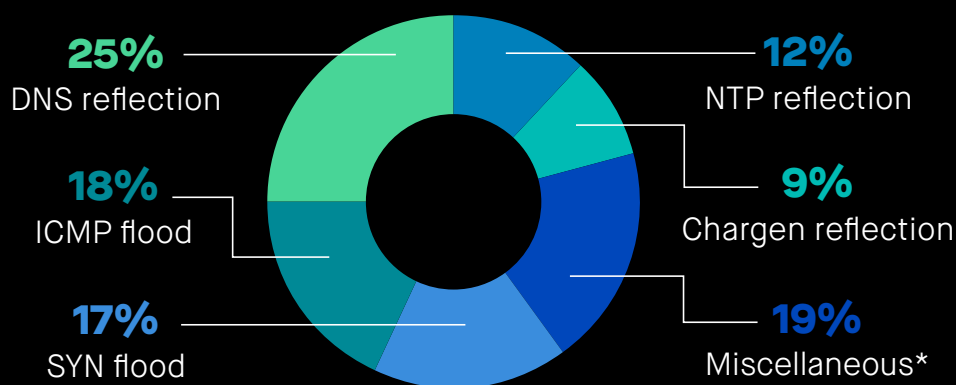
DDoS attack type breakdown:

Of the top 100 attacks, the majority employed three or four vectors within a single attack:

89% multivector

11% single vector

The percentages represent the dominant vector in the attacks mitigated.

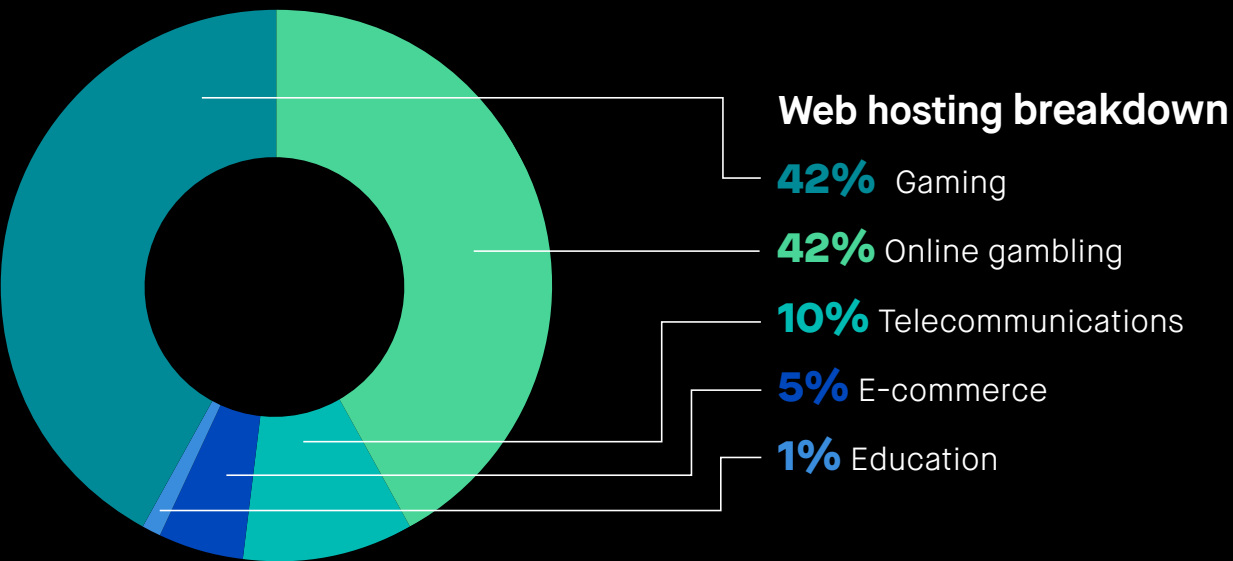
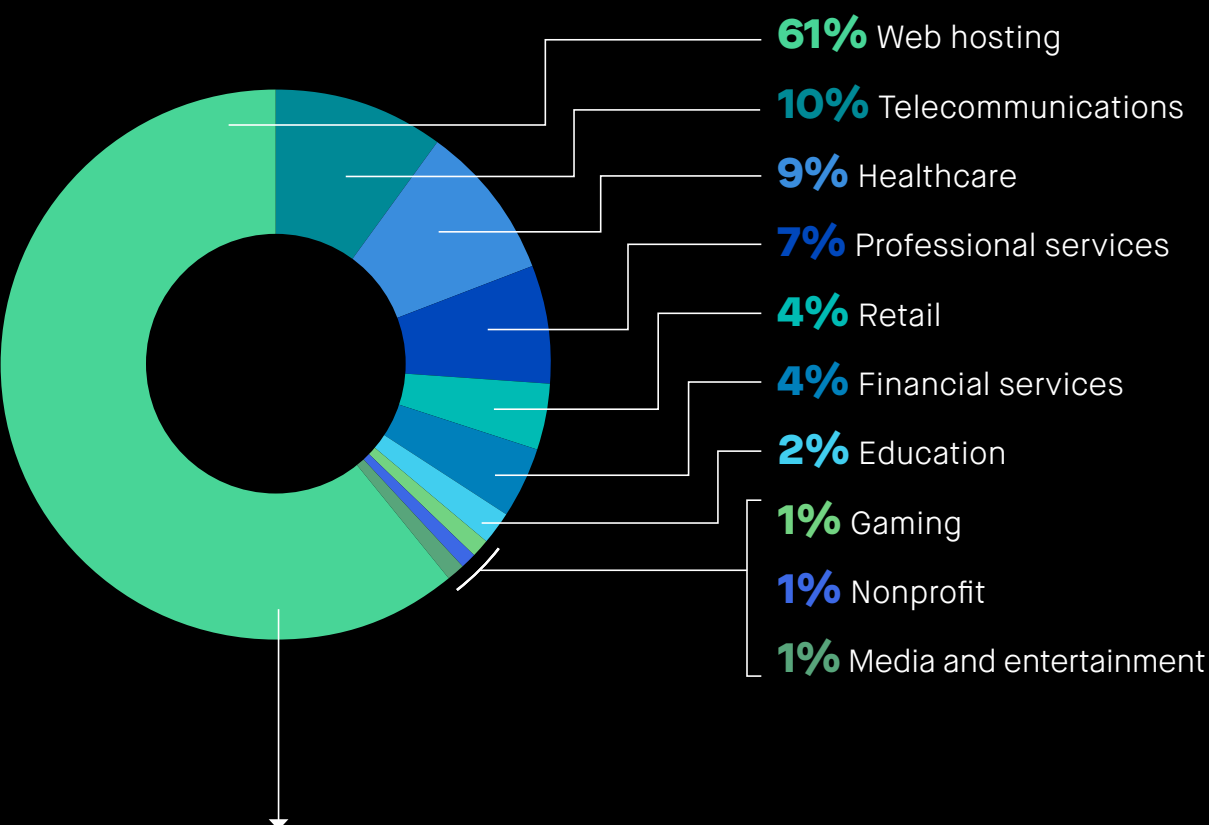


*HTTP, layer 7, UDP non-standard ports, etc.

Top 50 largest attacks by vector — 1H19

1. DNS reflection
2. ICMP flood - SYN flood: port 80 - SYN flood: port multiple
3. DNS reflection - ICMP flood - SYN flood: port multiple
4. Chargen reflection - DNS reflection - ICMP flood - NTP reflection - VPN SYN flood: port 30120 - SYN flood: port multiple
5. ICMP flood - NTP reflection
6. SYN flood: port 55903 - SYN flood: port 55919 - SYN flood: port 55970 - SYN flood: port multiple
7. DNS reflection
8. DNS reflection - ICMP flood
9. DNS reflection - ICMP flood - NTP reflection - SYN flood: port 30121 - SYN flood: port multiple
10. ICMP flood - NTP reflection
11. DNS reflection - ICMP flood - NTP reflection - SYN flood: port 28015 - SYN flood: port 30110
12. DNS reflection
13. DNS reflection - ICMP flood - SYN flood: port 25 - SYN flood: port multiple
14. DNS reflection - ICMP flood - NTP reflection - SSDP reflection - SYN flood: port 30120
15. SYN flood: port multiple
16. Chargen reflection
17. DNS reflection - ICMP flood - SYN flood: port 25
18. DNS reflection - ICMP flood
19. DNS reflection - ICMP flood
20. DNS reflection - ICMP flood - NTP reflection - RPC reflection - SSDP reflection - SYN flood: port 30120
21. ICMP flood - NTP reflection
22. Chargen reflection - ICMP flood - NTP reflection - SYN flood: port 443 - SYN flood: port multiple
23. ICMP flood - Chargen reflection
24. DNS reflection - ICMP flood - SYN flood: port 25 - SYN flood: port 53 - SYN flood: port 993 - SYN flood: port multiple
25. DNS reflection - ICMP flood - SYN flood: port multiple
26. DNS reflection - ICMP flood - SYN flood: port 5816
27. ICMP flood - Chargen reflection
28. DNS reflection - ICMP flood - NTP reflection - SYN flood: port 30122
29. DNS reflection - ICMP flood - SYN flood: port 25 - SYN flood: port multiple
30. DNS reflection
31. DNS reflection - ICMP flood
32. DNS reflection
33. DNS reflection - ICMP flood - NTP reflection - RPC reflection - SYN flood: port 30110 - SYN flood: port 30120 - SYN flood: port 30121 - SYN flood: port 3389 - SYN flood: port multiple
34. ICMP flood - NTP reflection
35. ICMP flood - NTP reflection - SSDP reflection - SYN flood: port 30120
36. ICMP flood - NTP reflection - SYN flood: port 30120
37. Chargen reflection - ICMP flood - NTP reflection - RPC reflection
38. DNS reflection - ICMP flood - RPC reflection - SYN flood: port 30120 - SYN flood: port multiple
39. ICMP flood - NTP reflection
40. ICMP flood - NTP reflection
41. ICMP flood - SYN flood: port 443
42. ICMP flood - NTP reflection - SSDP reflection
43. ICMP flood - NTP reflection - SYN flood: port 80 - SYN flood: port multiple
44. Chargen reflection - DNS reflection - ICMP flood - NTP reflection - RPC reflection
45. ICMP flood - NTP reflection - SYN flood: port 22 - SYN flood: port 443 - SYN flood: port multiple - UDP fragments
46. Chargen reflection - DNS reflection - ICMP flood - NTP reflection - SSDP reflection - SYN flood: port 443
47. Chargen reflection - DNS reflection - ICMP flood - NTP reflection - RPC reflection - SYN flood: port 28015 - SYN flood: port 29085
48. DNS reflection - ICMP flood - NTP reflection - SYN flood: port 22 - SYN flood: port 80 - SYN flood: port multiple
49. DNS reflection
50. DNS reflection - ICMP flood - NTP reflection

Top 500 largest DDoS attacks by industry — 1H19



Be bold — but beware

As these findings indicate, cyberthreats are escalating faster than many firms can identify, block and mitigate them. Visibility into the expanding threat landscape is imperative, but it's even more essential to act. As a security-minded community, we need continued collaboration to improve security policies, programs and platforms to meet the latest generation of threats.

What can you do to protect your business against current and future threats? Here are five recommendations:

- 1. Embed security into the network.** Networks are expanding faster than ever. The volume and velocity of network data is increasing — and will continue to skyrocket. Deploying better security at a network layer will help to provide a disproportional benefit due to the increased bandwidth that video, IoT, mobile and 5G will bring to the network.
- 2. Double down on the experts.** Security talent is scarce. Organizations must determine what is essential for them to architect and deliver internally and what can best be delivered by trusted partners. Think carefully about where you invest your top talent and where you can rely on others.
- 3. Make security simple.** When evaluating security solutions, keep these considerations top of mind: Does the solution improve your security, decrease your cost and reduce friction? Choose a security solution that meets at least two of these three criteria.
- 4. Close the security and engineering gap.** Security, IT and business engineering teams must work closely together to make sure that every product is secure. Security bolted onto the network late in the process can end up being disruptive to both the business and the user experience. It is essential that security be built into every product and solution as part of one seamless engineering process.
- 5. Change how we trust.** The days of implicit trust are over. A connected device tied to a person's identity cannot inherently be trusted. It's critical to look at applying security consistently across devices that are associated with a person's identity, regardless of who they are and where they are located.

Conclusion

Organizations will continue to accelerate the deployment of new technologies and capabilities that drive network traffic and security complexity. The explosion of data, edge computing and the increasing reliance on internet connectivity means that we must build security controls into the network layer to help protect the digital business.

To be usable, security must be simple. If security interrupts the ability of today's digital business to acquire, analyze and act upon data, its value is eroded. Protecting the digital business — including applications and data — means that security and the network must be connected to maximize data-driven value.

CenturyLink provides network-based layers of protection to help defend against an increasingly complicated threat landscape. We have one of the largest and most deeply peered IP backbones in the world, giving us expansive, near-real-time visibility into the threat landscape. Through our continued investment in Black Lotus Labs, we have harnessed the power of our global visibility to disrupt malicious actors.

Our direct actions aim to impact the ability of malicious actors to operate mass malware at scale and become too powerful. This means less malicious traffic hitting customer firewalls and entering internal networks, reducing the number of events that overburdened security teams must investigate. CenturyLink is focused on embedding security into our network services and integrating our global threat intelligence directly into our solutions.

As good stewards of the internet, we are proactively protecting global connectivity which benefits you whether you are a CenturyLink customer or not. We're doing this not for our bottom line, but for the greater good. We feel it's our unique responsibility as a global communications provider to act upon our vantage point and turn our backbone into a threat sensor and proactive defense platform. It means we are committed to doing our part to help keep the internet clean for everyone.

Want to protect your network? Learn more about [Connected Security](#) from CenturyLink. Contact an expert at **800-871-9244**.

If you would like to collaborate with [Black Lotus Labs](#) on threat research, please contact us on Twitter [@BlackLotusLabs](#).



This document is provided for informational purposes only and may require additional research and substantiation by the end user. In addition, the information is provided "as is" without any warranty or condition of any kind, either express or implied. Use of this information is at the end user's own risk. CenturyLink does not warrant that the information will meet the end user's requirements or that the implementation or usage of this information will result in the desired outcome of the end user. This document represents CenturyLink's products and offerings as of the date of issue.

Services not available everywhere. CenturyLink may change or cancel products and services or substitute similar products and services at its sole discretion without notice. ©2019 CenturyLink. All Rights Reserved.