Cisco 2018 Annual Cybersecurity Report

At-A-Glance: Cloud Security

In 2017, we saw a continued use of old tactics, such as phishing, and a rise in new types of attacks and techniques. On top of that, where users work, how they connect, and the applications they use continued to evolve. Our Cloud Security portfolio, which includes Umbrella and Cloudlock, can help address these risks and secure access for users, wherever they work. And with Cisco Umbrella Investigate, researchers and incident response teams can access threat intelligence about domains, IPs, and malware across the internet.

Whether it's phishing, spear phishing, or malicious tactics like domain squatting, Umbrella proactively blocks requests before a connection is established. So as malicious actors continue to use old techniques (or more advanced new ones), a user will be protected, even if they click a malicious link or attempt to connect to a malicious domain. Using Umbrella Investigate, our research team determined how often adversaries use, and reuse, registered-level domains (RLDs) in their attacks. We take these insights into known and emergent threats that are uncovered with Investigate and immediately block using Umbrella.

Organizations using cloud services need to be able to detect anomalous user behavior, including compromised accounts and malicious insiders and identify data exposures from over-sharing. Cloudlock uses machine-learning algorithms to provide a nuanced view of cloud user activity across cloud services. Cloudlock's User and Entity Behavior Analytics (UEBA) capabilities enable detection of anomalous behavior, including compromised accounts and malicious insiders. Using crowd-sourced security analytics, Cloudlock can detect and respond to connected cloud application (OAuth) ecosystem risks.

Cisco 2018 Security Capabilities Benchmark Study: Security viewed as a key benefit of hosting networks in the cloud

The use of on-premises and public cloud infrastructure is growing, according to the Cisco 2018 Security Capabilities Benchmark Study, although many organizations still host networks on-premises. In the 2017 study, 27 percent of security professionals said they are using off-premises private clouds, compared with 25 percent in 2016 and 20 percent in 2015 (Figure 1). Fifty-two percent said their networks are hosted on-premises as part of a private cloud.

Figure 1 More organizations are using private clouds

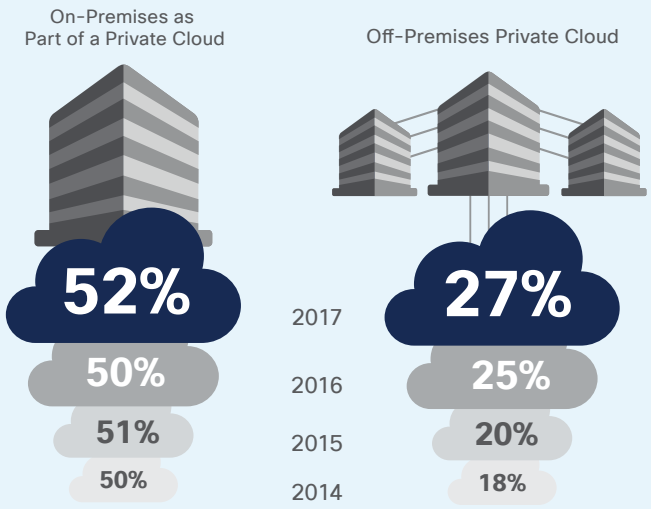

2014 (n=1727), 2015 (n=2417), 2016 (n=2887), 2017 (n=3625)

Source: Cisco 2018 Security Capabilities Benchmark Study

Of those organizations using the cloud, 36 percent host 25 to 49 percent of their infrastructure in the cloud, while 35 percent host 50 to 74 percent of their infrastructure in the cloud (Figure 2).

Security is the most common benefit of hosting networks in the cloud, according to the security personnel respondents. Among them, 57 percent said they host networks in the cloud because of better data security; 48 percent, because of scalability; and 46 percent, because of ease of use (see Figure 3).

Respondents also said that, as more infrastructure is moved to the cloud, they may look to invest in cloud access security brokers (CASBs) to add extra security to cloud environments.

Figure 2 Fifty-three percent of organizations host at least half of infrastructure in the cloud

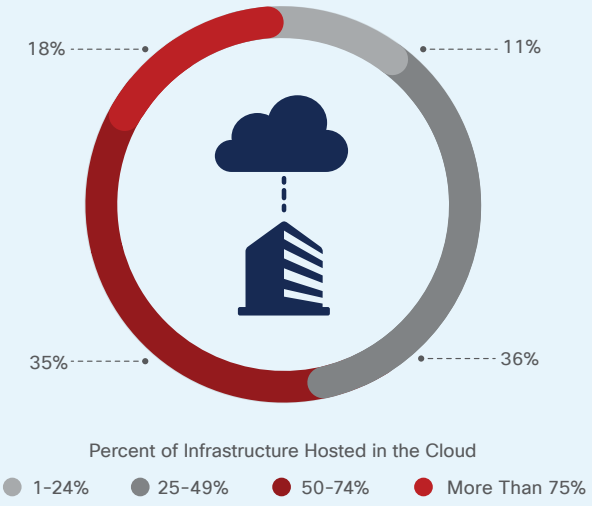

Source: Cisco 2018 Security Capabilities Benchmark Study

Figure 3 Fifty-seven percent believe the cloud offers better data security

Source: Cisco 2018 Security Capabilities Benchmark Study

Download the 2018 graphics at: cisco.com/go/acr2018graphics

ABUSE OF CLOUD SERVICES AND OTHER LEGITIMATE RESOURCES

As applications, data, and identities move to the cloud, security teams must manage the risk involved with losing control of the traditional network perimeter. Attackers are taking advantage of the fact that security teams are having difficulty defending evolving and expanding cloud and IoT environments. One reason is the lack of clarity around who exactly is responsible for protecting those environments.

To meet this challenge, enterprises may need to apply a combination of best practices, advanced security technologies like machine learning, and even some experimental methodologies, depending on the services they use for their business and how threats in this space evolve.

Malicious use of legitimate resources for backdoor C2

When threat actors use legitimate services for command and control (C2), malware network traffic becomes nearly impossible for security teams to identify because it mimics the behavior of legitimate network traffic. Adversaries have a lot of Internet "noise" to use as cover because so many people today rely on services like Google Docs and Dropbox to do their work, regardless of whether these services are offered or systemically endorsed by their employers.

Figure 4 shows several of the well-known legitimate services that researchers with Anomali, a Cisco partner and threat intelligence provider, have observed being used in malware backdoor C2 schemas[1] in the last few years. (Note: These types of services face a dilemma in combating abuse, as making it more difficult for users to set up accounts and use their services can adversely affect their ability to generate revenue.)

Figure 4 Examples of legitimate services abused by malware for C2

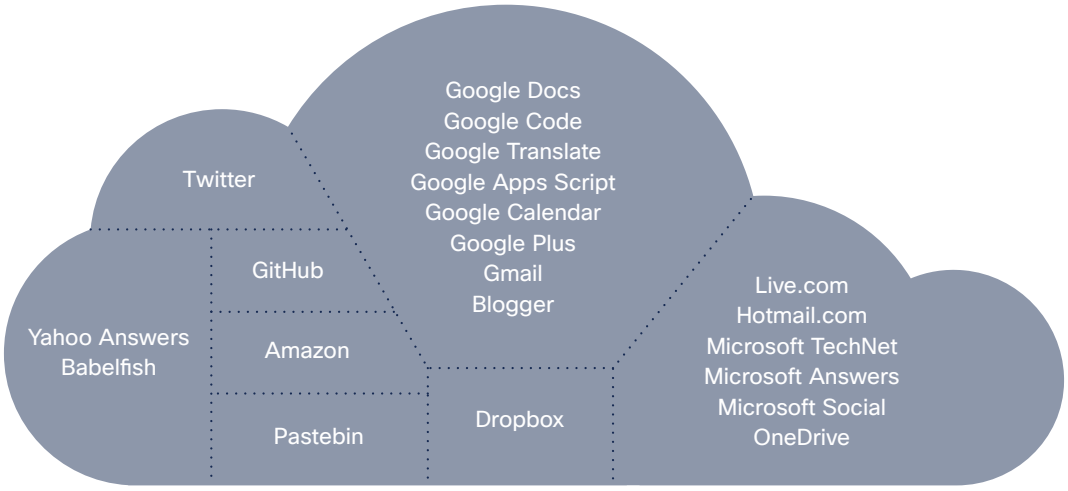

Source: Anomali

1 Anomali defines a C2 schema as "the totality of IP addresses, domains, legitimate services, and all the remote systems that are part of the ... communications architecture" of malware.

Extracting optimal value from resources

Cisco security researchers analyzed newly seen unique query names (domains) associated with DNS queries made over a seven-day period in August 2017. Note that "newly seen" in this discussion has no bearing on when a domain was created; it relates to when a domain was first "seen" by Cisco cloud security technology during the period of observation.

The purpose of this research was to gain more insight into how often adversaries use, and reuse, registered-level domains (RLDs) in their attacks. Understanding threat actor behavior at the domain level can help defenders identify malicious domains, and related subdomains, that should be blocked with first-line-of-defense tools like cloud security platforms.

So that our researchers could focus solely on the core group of unique RLDs—about 4 million in total—subdomains were stripped from the sample of newly seen domains. Only a small percentage of the RLDs in that sample were categorized as malicious. Of the RLDs that were malicious, more than half (about 58 percent) were reused, as Figure 5 shows.

Figure 5 Percent of new vs. reused domains

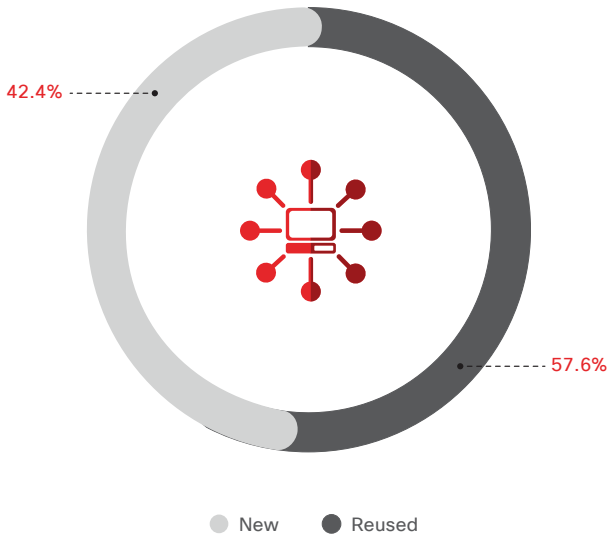

Source: Cisco Security Research

That finding suggests that, while most attackers build new domains for their campaigns, many are focused on trying to get the best return on their investments by launching multiple campaigns from a single domain. Domain registration can be costly, especially at the scale most attackers require to execute their campaigns and evade detection.

One-fifth of malicious domains quickly put into use

Adversaries may sit on domains for days, months, or even years after registering them, waiting for the right time to use them. However, Cisco threat researchers observed that a significant percentage of malicious domains—about 20 percent—were used in campaigns less than one week after they were registered (see Figure 6).

Figure 6 RLD registration times

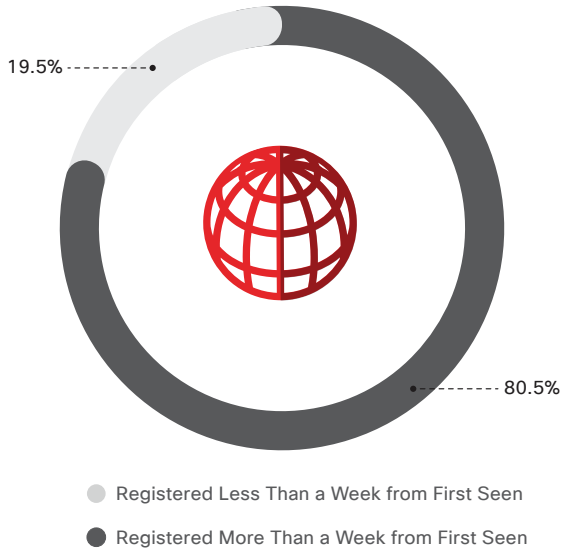

Source: Cisco Security Research

Many new domains tied to malvertising campaigns

Most malicious domains we analyzed were associated with spam campaigns—about 60 percent. Nearly one-fifth of the domains were connected to malvertising campaigns (see Figure 7). Malvertising has become an essential tool for directing users to exploit kits, including those that distribute ransomware.

Figure 7 Malicious categorizations

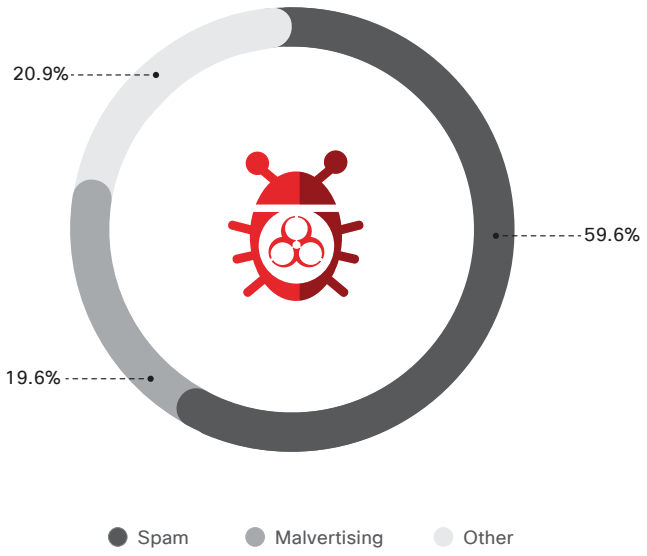

Source: Cisco Security Research

Well-worn, domain-related techniques for creating malvertising campaigns include domain shadowing. In this technique, attackers steal legitimate domain account credentials to create subdomains directed at malicious servers. Another tactic is the abuse of free, dynamic DNS services to generate malicious domains and subdomains. That allows threat actors to deliver malicious payloads from constantly changing hosting IPs, either infected users' computers or compromised public websites.

Domains reuse infrastructure resources

The malicious RLDs in our sample also appeared to reuse infrastructure resources, such as registrant email addresses, IP addresses, autonomous system numbers (ASNs), and nameservers (see Figure 8). This is further evidence of adversaries trying to get the most value from their investments in new domains and preserve resources, according to our researchers.

For example, an IP address can be used by more than one domain. So, an attacker laying the groundwork for a campaign might decide to invest in a few IP addresses and an array of domain names instead of servers, which cost more.

Figure 8 Reuse of infrastructure by malicious RLDs

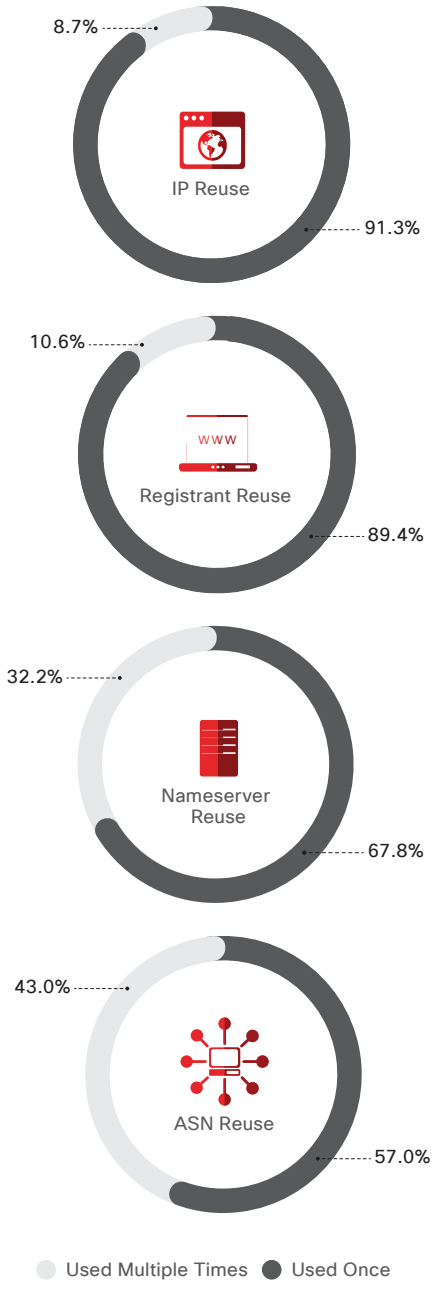

Source: Cisco Security Research

The resources that RLDs reuse give clues to whether the domain is likely to be malicious. For example, reuse of registrant emails or IP addresses occurs infrequently, so a pattern of reuse on either front suggests suspicious behavior. Defenders can have a high degree of confidence in blocking those domains, knowing that doing so probably will have no negative impact on business activity.

Static blocking of ASNs and nameservers is not likely to be feasible in most cases. However, patterns of reuse by RLDs are worthy of further investigation to determine whether certain domains should be blocked.

Using intelligent, first-line-of-defense cloud security tools to identify and analyze potentially malicious domains and subdomains can help security teams follow the trail of an attacker and answer questions, such as:

- What IP address does the domain resolve to?
- What ASN is associated with that IP address?
- Who registered the domain?
- What other domains are associated with that domain?

The answers can help defenders not only refine security policies and block attacks, but also prevent users from connecting to malicious destinations on the Internet while they're on the enterprise network.

Insider threats: Taking advantage of the cloud

In previous security reports, we have discussed the value of OAuth permissions and super-user privileges to enforce who can enter networks, and how they can access data.[1] To further examine the impact of user activity on security, Cisco threat researchers recently examined data exfiltration trends. They employed a machine-learning algorithm to profile 150,000 users in 34 countries, all using cloud service providers, from January to June 2017. The algorithm accounted for not only the volume of documents being downloaded, but also variables such as the time of day of downloads, IP addresses, and locations.

After profiling users for six months, our researchers spent 1.5 months studying abnormalities. flagging 0.5 percent of users for suspicious downloads. That's a small amount, but these users downloaded, in total, more than 3.9 million documents from corporate cloud systems, or an average of 5200 documents per user during the 1.5-month period. Of the suspicious downloads, 62 percent occurred outside of normal work hours; 40 percent took place on weekends.

Cisco researchers also conducted a text-mining analysis on the titles of the 3.9 million suspiciously downloaded documents.

One of the most popular keywords in the documents' titles was "data." The keywords most commonly appearing with the word "data" were "employee" and "customer." Of the types of documents downloaded, 34 percent were PDFs and 31 percent were Microsoft Office documents (see Figure 9).

Applying machine-learning algorithms offers a more nuanced view of cloud user activity beyond just the number of downloads. In our analysis, 23 percent of the users we studied were flagged more than three times for suspicious downloads, usually starting with small numbers of documents. The volume slowly increased each time, and eventually, these users showed sudden and significant spikes in downloads (Figure 10).

Machine-learning algorithms hold the promise of providing greater visibility into the cloud and user behavior. If defenders can start predicting user behavior in terms of downloads, they can save the time it might take to investigate legitimate behavior. They can also step in to stop a potential attack or data-exfiltration incident before it happens.

Figure 9 Most commonly downloaded documents

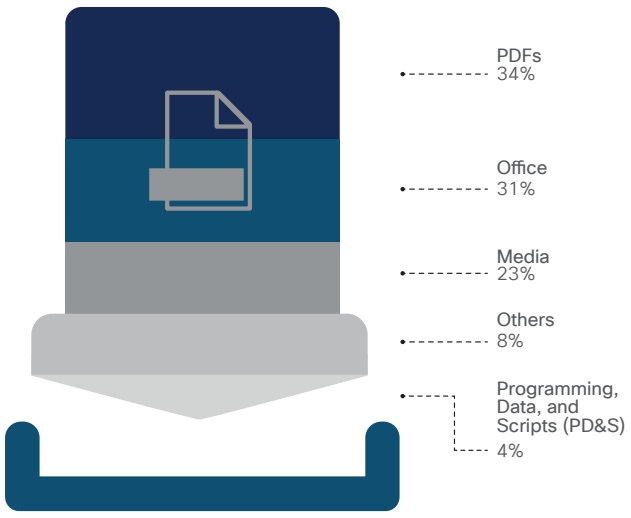

Source: Cisco Security Research

Figure 10 Machine-learning algorithms capture suspicious user download behavior

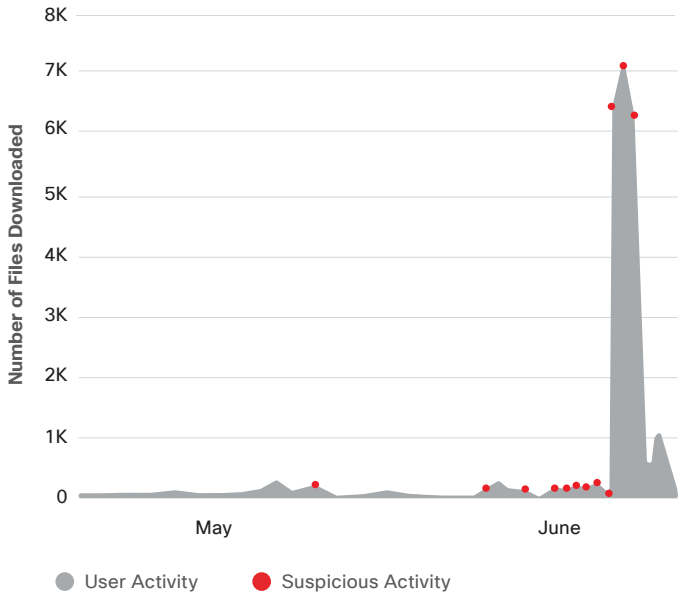

Source: Cisco Security Research

1 *Cisco 2017 Midyear Cybersecurity Report*: cisco.com/c/m/en_au/products/security/offers/cybersecurity-reports.html.

EMAIL THREATS

No matter how much the threat landscape changes, malicious email and spam remain vital tools for adversaries to distribute malware because they take threats straight to the endpoint. By applying the right mix of social engineering techniques, such as phishing and malicious links and attachments, adversaries need only to sit back and wait for unsuspecting users to activate their exploits.

Social engineering still a critical launchpad for email attacks

Phishing and spear phishing are well-worn tactics for stealing users' credentials and other sensitive information, and that's because they are very effective. In fact, phishing and spear phishing emails were at the root of some of the biggest, headline-grabbing breaches in recent years. Two examples from 2017 include a widespread attack that targeted Gmail users[1] and a hack of Irish energy systems.[2]

To gauge how prevalent phishing URLs and domains are on today's Internet, Cisco threat researchers examined data from sources that investigate potentially "phishy" emails submitted by users through community-based, anti-phishing threat intelligence. Figure 11 shows the number of phishing URLs and phishing domains observed during the period from January to October 2017.

Figure 11 Number of observed phishing URLs and domains by month

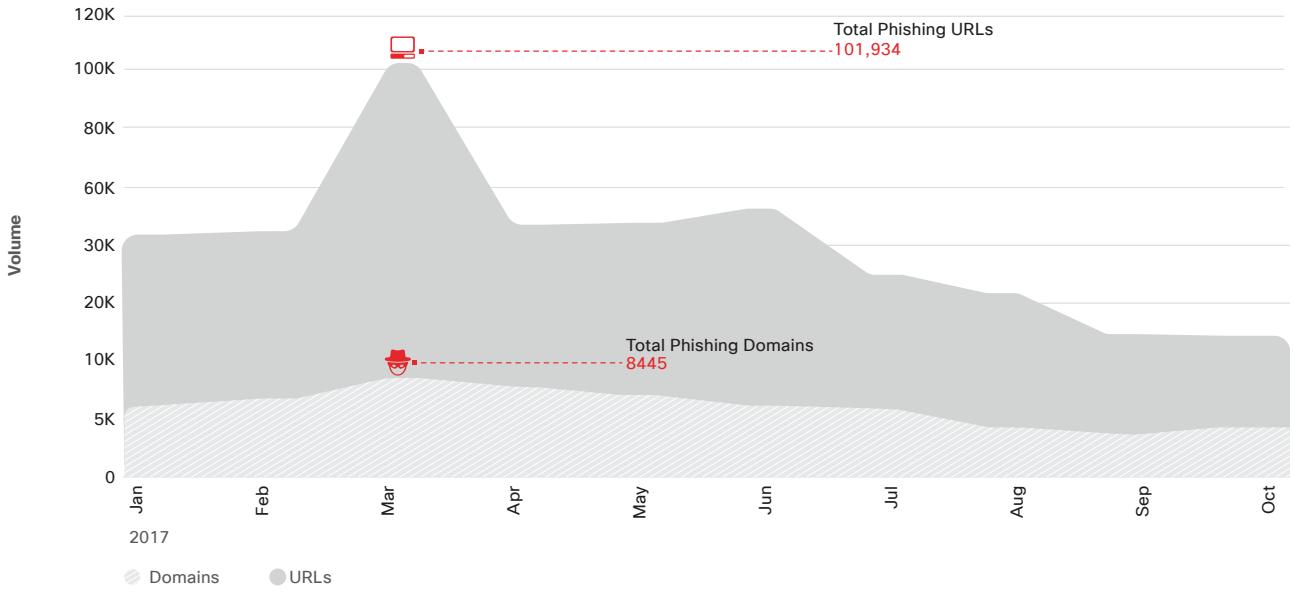

Source: Cisco Security Research

The spikes seen in March and June can be attributed to two different campaigns. The first appeared to target users of a major telecom services provider. That campaign:

- Involved 59,651 URLs containing subdomains under aaaainfomation[dot]org.
- Had subdomains that contained random strings consisting of 50-62 letters.

Each subdomain length (50-62) contained about 3500 URLs, which allowed for programmatic use of the subdomains (example: Cewekonuxykysowegulukozapojygepuqybyteqe johofopefogu[dot]aaaainfomation[dot]org).

Adversaries used an inexpensive privacy service to register the domains observed in this campaign.

1 *Massive Phishing Attack Targets Gmail Users,* by Alex Johnson, NBC News, May 2017: nbcnews.com/tech/security/massive-phishing-attack-targets-millions-gmail-users-n754501.
2 *Hackers target Irish energy networks amid fears of further cyber attacks on UK's crucial infrastructure,* by Lizzie Deardon, *The Independent*, July 2017: independent.co.uk/news/world/europe/cyber-attacks-uk-hackers-target-irish-energy-network-russia-putin-electricity-supply-board-nuclear-a7843086.html.

In the second campaign, which was most active in June, threat actors used the name of a legitimate tax agency in the United Kingdom to disguise their actions. They employed 12 top-level domains (TLDs). Eleven of the domains were URLs with six random six-character strings (example: jyzwyp[dot]top). And nine of the domains were associated with more than 1600 phishing sites each.

Like the March campaign, adversaries registered the domains using a privacy service to conceal domain registration information. They registered all the domains over a two-day period. On the second day, nearly 19,000 URLs connected to the campaign were observed, and all were discovered within a five-hour window (for more on how quickly threat actors put newly registered domains to use, see "Malicious use of legitimate resources for backdoor C2," on page 3).

TLD distribution across known phishing sites

Our analysis of phishing sites during the period from January to August 2017 found that threat actors were employing 326 unique TLDs for these activities, including .com, .org, .top (largely due to the United Kingdom taxing agency campaign), and country-specific TLDs (see Figure 12). Employing lesser-known TLDs can be advantageous for adversaries; these domains are typically inexpensive and often offer inexpensive privacy protection.

Figure 12 TLD distribution across known phishing sites

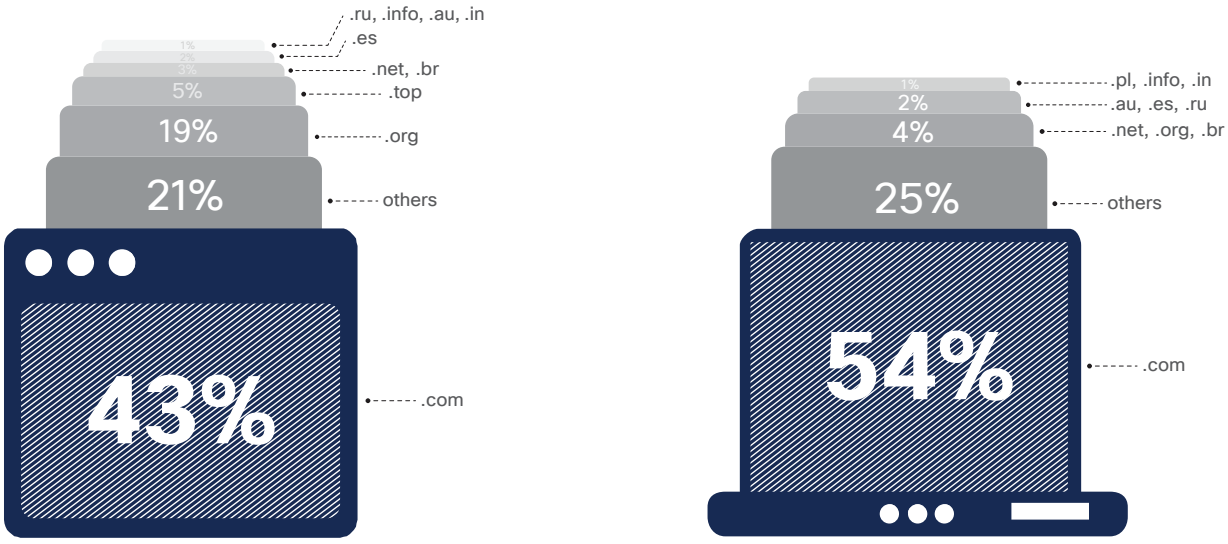

Source: Cisco Security Research

Defenders should be vigilant in monitoring this "old" threat

In 2017, tens of thousands of phishing attempts were reported monthly to the community-based, anti-phishing threat intelligence services included in our analysis. Some of the common tactics and tools adversaries use to execute phishing campaigns include:

- **Domain squatting:** Domains named to look like valid domains (example: cisc0[dot]com).
- **Domain shadowing:** Subdomains added under a valid domain without the owner's knowledge (example: badstuff[dot]cisco[dot]com).
- **Maliciously registered domains:** A domain created to serve malicious purposes (example: viqpbe[dot]top).
- **URL shorteners:** A malicious URL disguised with a URL shortener (example: bitly[dot]com/random-string). Note: In the data we examined, Bitly.com was the URL-shortening tool adversaries used most. Malicious shortened URLs represented 2 percent of the phishing sites in our study. That number peaked to 3.1 percent in August.
- **Subdomain services:** A site created under a subdomain server (example: mybadpage[dot]000webhost[dot]com).

Threat actors in the phishing and spear phishing game are continuously refining social engineering methods to trick users into clicking malicious links or visiting fraudulent web pages, and providing credentials or other types of high-value information. User training and accountability, and the application of email security technologies, remain crucial strategies for combating these threats.

Recommendations for defenders

When adversaries inevitably strike their organizations, will defenders be prepared, and how quickly can they recover? Findings from the **Cisco 2018 Security Capabilities Benchmark Study**—which offers insights on security practices from more than 3600 respondents across 26 countries—show that defenders have a lot of challenges to overcome.

Even so, defenders will find that making strategic security improvements and adhering to common best practices can reduce exposure to emerging risks, slow attackers' progress, and provide more visibility into the threat landscape. They should consider:

- Implementing first-line-of-defense tools that can scale, like cloud security platforms.
- Confirming that they adhere to corporate policies and practices for application, system, and appliance patching.
- Employing network segmentation to help reduce outbreak exposures.
- Adopting next-generation endpoint process monitoring tools.
- Accessing timely, accurate threat intelligence data and processes that allow for that data to be incorporated into security monitoring and eventing.
- Performing deeper and more advanced analytics.
- Reviewing and practicing security response procedures.
- Backing up data often and testing restoration procedures—processes that are critical in a world of fast-moving, network-based ransomware worms and destructive cyber weapons.
- Reviewing third-party efficacy testing of security technologies to help reduce the risk of supply chain attacks.
- Conducting security scanning of microservice, cloud service, and application administration systems.
- Reviewing security systems and exploring the use of SSL analytics—and, if possible, SSL decryption.

Defenders should also consider adopting advanced security technologies that include machine learning and artificial intelligence capabilities. With malware hiding its communication inside of encrypted web traffic, and rogue insiders sending sensitive data through corporate cloud systems, security teams need effective tools to prevent or detect the use of encryption for concealing malicious activity.

WANT MORE INFORMATION?

The full 2018 Annual Cybersecurity Report is just a click away. Get full report.