RAPID7

RAPID7 RESEARCH

QUARTERLY THREAT REPORT

By Rebekah Brown, Threat Intelligence Lead, Rapid7, Inc.
Kwan Lin, Senior Data Scientist, Rapid7, Inc.
Bob Rudis, Chief Data Scientist, Rapid7, Inc.

May 15, 2018

Q1 2018

CONTENTS

WHAT IS A THREAT?

We throw the term “threat” around a lot, and so it’s important to define exactly what it is we mean.

When there is an adversary with the intent, capability, and opportunity, a **threat** exists.

When two or more of these elements are present (e.g. intent and capability, but no opportunity), we call it an **impending threat**, because there is just one missing piece before it becomes a true threat.

When there is just one element present (e.g. an opportunity in the form of a software vulnerability), we call it a **potential threat**. There is the potential for it to turn into a true threat, although there are additional components that need to come to fruition before it has a real impact to most organizations.

INTRODUCTION

Spring is here! The sun is shining, the birds are chirping, and attackers are coming up with more convincing ways to steal user credentials. While fairer weather does not lull attackers into slowing their pace, it does mean that you can at least sit in the sunshine and read our findings from the past quarter before continuing the mission of defending your network against an often persistent, sometimes creative, and always-on-the-job adversary.

This quarter's report covers three main areas of concern for the modern IT defender:

- First, credential theft, reuse, and subsequent suspicious logins are—today—the most commonly reported significant incident we're seeing across both small (<1,000 endpoints) and large organizations (≥1,000 endpoints).
- Second, the DDoS landscape just got a lot more interesting with the debut of a new technique using misconfigured—and plentiful—memcached servers.
- Finally, we take a look at the increasing levels of SMB and Cisco SMI attacker probes and attacks, where the former continues to define the "new normal" level of background malicious behavior around Windows networking, and the latter begins to bring shape to this relatively new attack vector targeting core router infrastructure.

What follows is a breakdown of trends we saw throughout 2018 Q1, including what we're referring to as "significant investigations," takeaways for the next quarter, and an overview of our methodologies and the resources at our disposal when crafting this report.

Figure 1: Q1 Threat Event Distribution by Industry

Normalized by number of events per organization per industry for Q1 2018. Columns sum to 100% in-industry

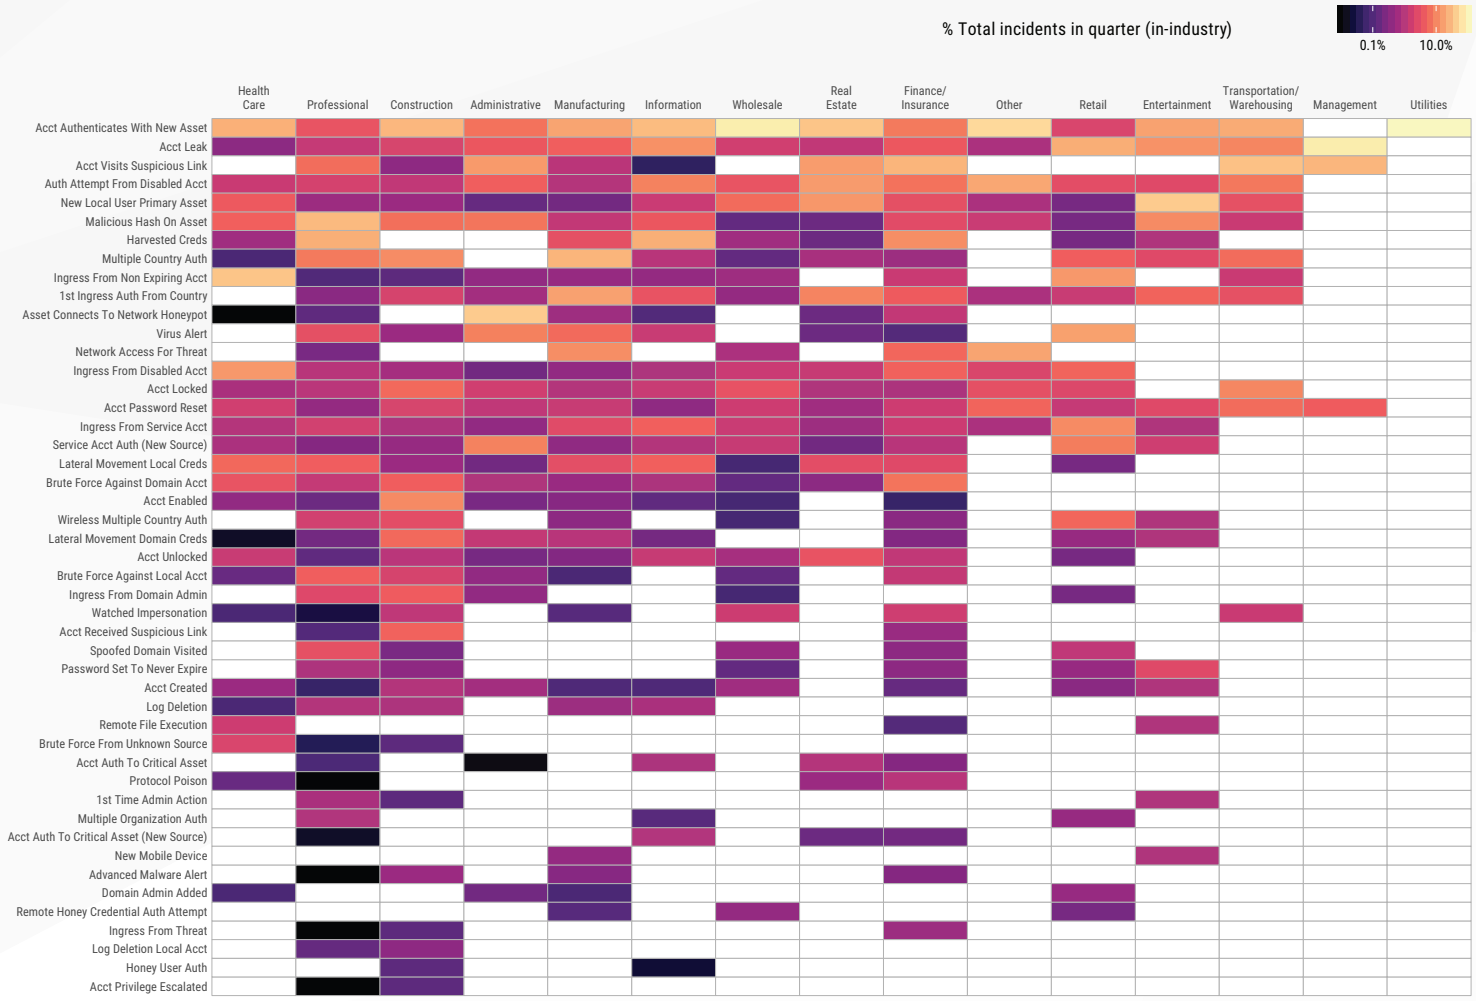

Source: Rapid7 Managed Detection and Response

Targeting Healthcare

Our industry snapshot for the first quarter of 2018 shows a continuing trend away from a primary focus on financial, professional, and administrative industries as adversaries look toward other industries with valuable data. In our 2017 wrap-up we highlighted increases in targeting activity in the real estate and construction industries, and this past quarter we saw a large increase in activity against the healthcare sector, so much so that it is our top-targeted industry for 2018 Q1.

The healthcare sector has been a desirable target for adversaries for some time, with attacks increasing since at least 2015. It is also an industry where security has even higher stakes; healthcare organizations often have a complex, distributed IT infrastructure with difficult-to-patch legacy systems and proprietary medical devices, making them challenging to secure quickly. They also rely on system availability to keep operations running when lives are on the line, and adversaries have frequently targeted that availability using tactics such as ransomware or telephonic denial of service attacks (TDoS) to overwhelm critical phone lines.

Healthcare also holds a great deal of sensitive data, both financial and personally identifiable information, that attackers have clearly shown they are interested in stealing. Symantec's report on the Orangeworm attack group details espionage-type activity targeting the healthcare sector dating back to at least 2015. From our data set, we have identified that a large number of the attack vectors include remote access, including suspicious logins, access attempts from disabled accounts, and account leaks. With the seemingly unending account and credential leaks, it is important to use two-factor authentication whenever possible and to identify and remediate instances where an employee's credentials may have been compromised to ensure that adversaries cannot use them to access the networks.

In addition to healthcare, we continued to see increases in activity against construction, manufacturing, and wholesale business operation. In most cases these industries still do not experience the same breadth of activity as we see against the financial and professional sectors, indicating that adversaries are focusing heavily on a few attack vectors against these industries.

Increasing Dangerous User Behavior

Figure 2: Q1 Incident Frequency by Organization Size

Difference (% of events per group) between non-hash-based detected threat events by organization size

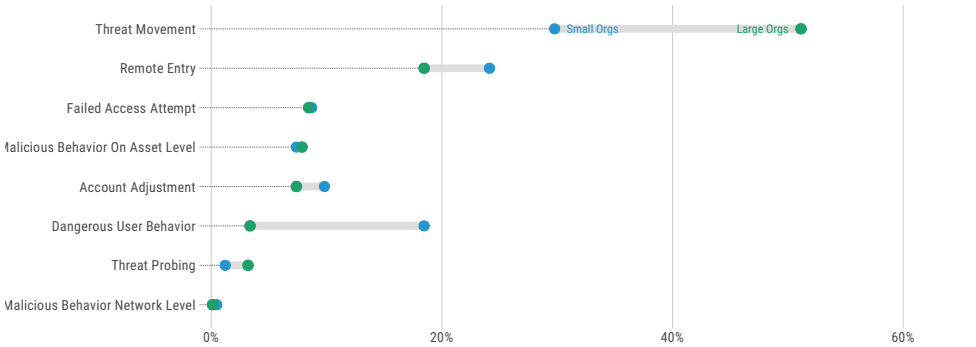

'Large Orgs' defined as >= 1,000 MDR monitored assets
Source: Rapid7 Managed Detection and Response

Incident frequency by organization size also had some surprises this quarter. Although threat movement and remote entry remained the top incident types (continuing the trends we identified in 2017), this quarter we saw a large increase in dangerous user behavior, including users visiting malicious sites or installing and running questionable programs. This jump was most notable in large organizations: where we saw 12% of dangerous user behavior incidents in 2017, they account for nearly 35% of incidents in the first quarter of 2018.

While dangerous user behavior increased this quarter across all organizations, remote entry attempts went down for larger organizations and increased for smaller organizations. In our Quarterly Threat Report wrap-up for 2017, we identified that large organizations often have a bigger security staff focused on identifying and limiting exposed systems. In the first quarter of 2018 we had several vulnerabilities disclosed that allowed remote access, including Cisco Smart Install; several campaigns targeting exposed systems, such as GoScanSSH; and the continued use of EternalBlue in ransomware campaigns. It is critical that all organizations, both large and small, identify exposed systems to ensure that they are up to date on patches and close any ports unnecessary for normal production activities.

Adversaries can also gain remote access by using legitimate credentials to log in to third-party services such as Dropbox, Office365, DocuSign, and a variety of other services that organizations regularly leverage for business operations. These credentials can be exposed in several different ways: they can be guessed or brute-forced, they can be disclosed as part of a large credential account leak, or they can be stolen from a user through phishing campaigns designed to trick users into entering their credentials. In Q1 we saw that the most common phishing campaigns were aimed at stealing credentials and masqueraded as DocuSign, Office365, and Dropbox, although there were other attempts to masquerade as Amazon Prime, Apple, and other sites or services.

Adversaries can also gain remote access by using legitimate credentials to log in to third-party services such as Dropbox, Office365, DocuSign, and a variety of other services that organizations regularly leverage for business operations.

Once credentials are obtained, an adversary could then log in to the service as if they were the user. In some cases the login activity would be flagged as malicious if it comes from a suspicious location, such as foreign countries that the organization does not do business in, but stealthy adversaries will find ways to blend in with normal login traffic, making it harder to identify their activity. By and large, adversaries attempted to steal credentials from employees at larger organizations at nearly three times the rate of smaller organizations, but much of that is likely due to the number of employees available to target.

Figure 3: Weekly Significant Investigations/Incidents Involving Credential Theft/Usage

Ordered by most frequency credential incident type
Percentage is the fraction of all incidents that week (not just credential-oriented incidents)

	Week (ending — Sunday)												
	Jan-07	Jan-14	Jan-21	Jan-28	Feb-04	Feb-11	Feb-18	Feb-25	Mar-04	Mar-11	Mar-18	Mar-25	Apr-01
Docusign Page		5.6%		2.5%	3.8%	10.5%	5.0%		3.6%	4.3%		6.7%	4.2%
Office 365 Page	18.2%	5.6%		5.0%			5.0%			4.3%		3.3%	12.5%
Dropbox Page	18.2%		8.3%	2.5%					3.6%				4.2%
System Recovery		11.1%		2.5%	3.8%								
One Drive										4.3%			4.2%
G Suite Login Page												6.7%	
Amazon Prime Page				2.5%									
Apple Id Page										2.2%			
At&T Page						5.3%							
Google Doc Page						5.3%							

Source: Rapid7 Managed Detection & Response

Aside from the two incident types listed above, incident frequency was fairly close between large and small organizations, and all were fairly low frequency overall. While it's unlikely these trends will stay the same forever (nothing ever does), it does give organizations, both large and small, something to focus on in the near term.

Figure 4: Credential Threat Actions by Org Size

When looking at all the credential-related threat events, most in our corpus are headed towards large organizations

Source: Rapid7 Managed Detection & Response

SIGNIFICANT INVESTIGATIONS

While we like to look at trends in alerts over time, there is almost never a one-alert-per-incident correlation. Adversary actions involve multiple steps, which generate multiple alerts, and after analysis, tell the story of what actually happened in the incident. This year, in addition to reviewing alert trends, we are also capturing the stories that those alerts tell, which we call "significant investigations."

Figure 5: Weekly Significant Investigations/Incidents Counts

Weekly percentage is the percentage of incidents that week out of all incidents in Q1 2018

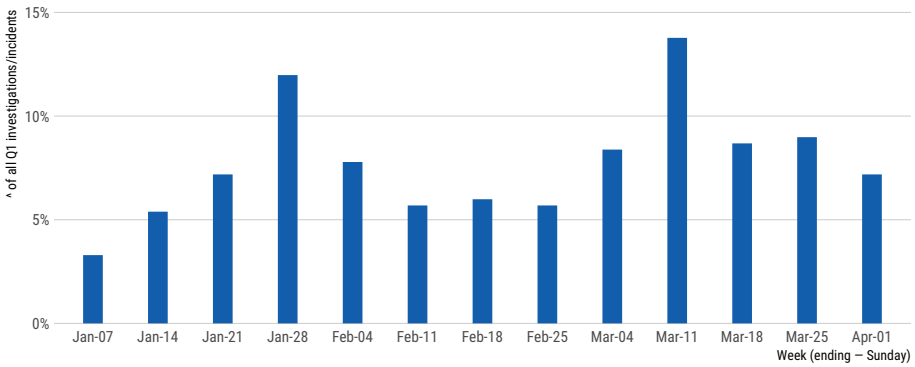

Source: Rapid7 Managed Detection & Response

What we can see from a weekly view of significant investigations is that it is hard to identify a "normal" week in reporting. Some weeks, such as those around holidays at the beginning of the year and during February, have fewer significant incidents. However, there are also weeks of heavy activity resulting in higher-than-normal reporting. Two of these heavy weeks were the weeks of Jan 28 and March 11. In both cases, these higher reporting weeks involved heavier-than-normal phishing and subsequent suspicious login activity, rather than any high-profile attacks or vulnerabilities in the news during those weeks. While it is important to maintain an awareness of significant attacks and vulnerabilities being reported, it is critical to stay vigilant for traditional malicious activity, because these activities continue regardless of whether there is something more interesting going on in the news.

Learn more about how we turn alert trends into Attacker Behavior Analytics. ▶▶

The top four significant incident types in this quarter were suspicious logins, phishing, malware on system, and cryptocurrency mining. The significant amount of suspicious logins correlates to the large number of remote entry alerts identified throughout the quarter, and also ties in to the second-highest threat identified: phishing. The majority of phishing in Q1 of 2018 involved sending a user to sites mimicking authentication sites that are designed to steal credentials, subsequently enabling attackers to log in to the network.

What is more, these top four attacks were seen at a much higher rate in larger organizations than in smaller organizations, whereas the majority of other attacks were seen at nearly identical rates between large and small organizations. The most striking contrasts were with suspicious logins, phishing, and malware. Suspicious logins accounted for roughly 75% of in-category significant incidents in large organizations and 23% in small; phishing accounted for 80% in large organizations and 20% in small; and malware on the system accounted for 85% of in-category significant incidents in large organizations and 20% in small.

Figure 6: Threat Event Summary

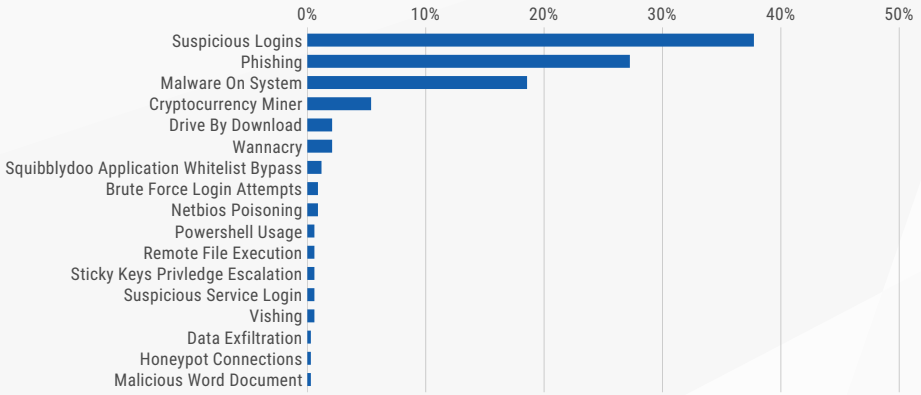

Source: Rapid7 Managed Detection & Response

Figure 7: Significant Incident/Investigation Threat Events By Org Size

Percentage is the ratio of the incident threat event count by total incidents

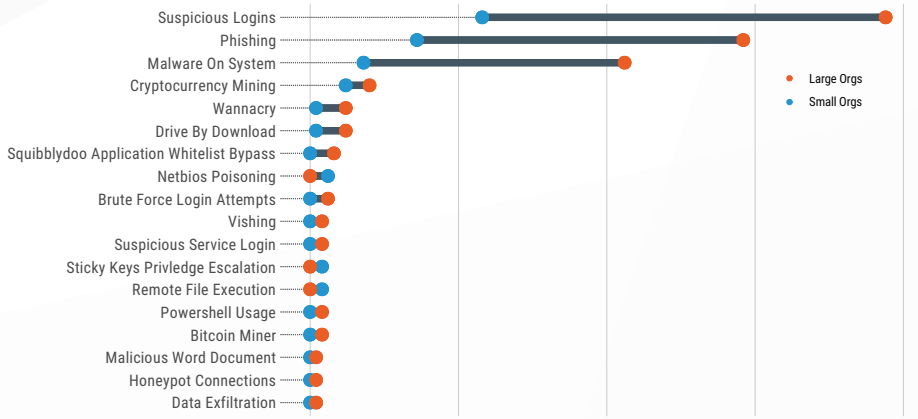

Source: Rapid7 Managed Detection & Response

TAKING INVENTORY: AMPLIFICATION DDoS ATTACKS IN 2018 Q1

In March, US-CERT updated an existing **alert document** on "UDP-Based Amplification Attacks" and added **memcached**—a free and open source, high-performance, distributed memory object caching system—to the list of known/active amplification DDoS protocols after a record-breaking 30-plus minute amplification DDoS attack on GitHub that topped out at over 1.35Tbps.

Amplification attacks work when an attacker sends a command or protocol exchange to an unwitting service using a spoofed source address. The service then sends the reply to that spoofed address instead of the system(s) the attacker is using. By creating and sending thousands and thousands of these spoofed requests, attackers flood a target system or network with unsolicited replies. Attackers can gain a significant attack benefit by choosing a service that sends a large reply using only a small command or protocol exchange. This is called the bandwidth amplification factor (BAF). A BAF of 1 means that for every byte sent by the attacker to the vulnerable service, said service responds with 1 byte to the spoofed target.

Figure 8: Maximum Amplification Attack Capacity for Known Amplification-DDoS Susceptible Protocols

With thousands of active and vulnerable internet hosts to choose from, Memcached — the new kid on the DDoS block — will likely be used in future large-scale attacks

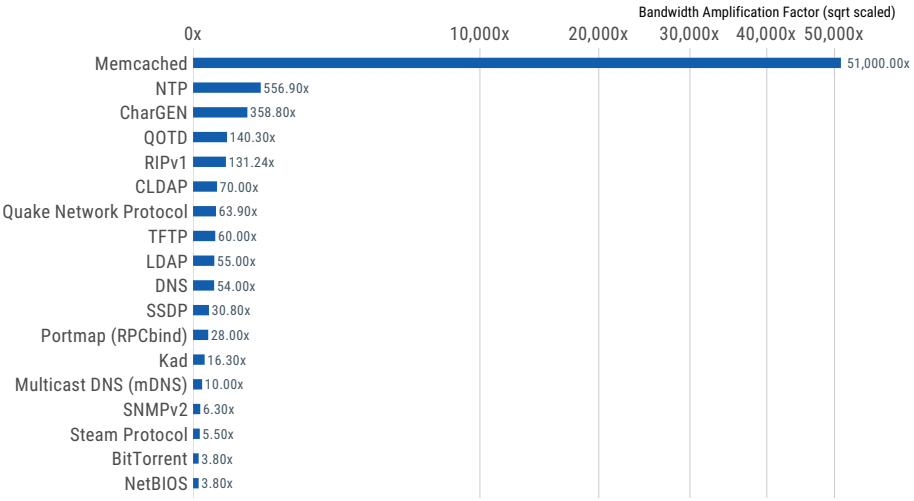

Source: US-CERT <https://www.us-cert.gov/ncas/alerts/TA14-017A>

Up until memcached, NTP had the top spot with a BAF of over 550. NTP is now in a distant second place since memcached has a maximum BAF potential of 51,000 and turns out to be an especially handy attack tool since there are tens of thousands of open, vulnerable instances on the internet.

Figure 9: 'Inventory' Scans for Hosts That Can Be Used in DDoS Amplification (DDoS-A) Attacks

A number of organizations suffered through DDoS-A attacks in Q1 2018.

Prior to these attacks, Rapid7 researchers captured evidence of adversaries 'taking inventory' of hosts on the internet capable of being used in these attacks.

The chart shows the most popular, current DDoS-A services with clear spikes when the inventory scans are taking place.

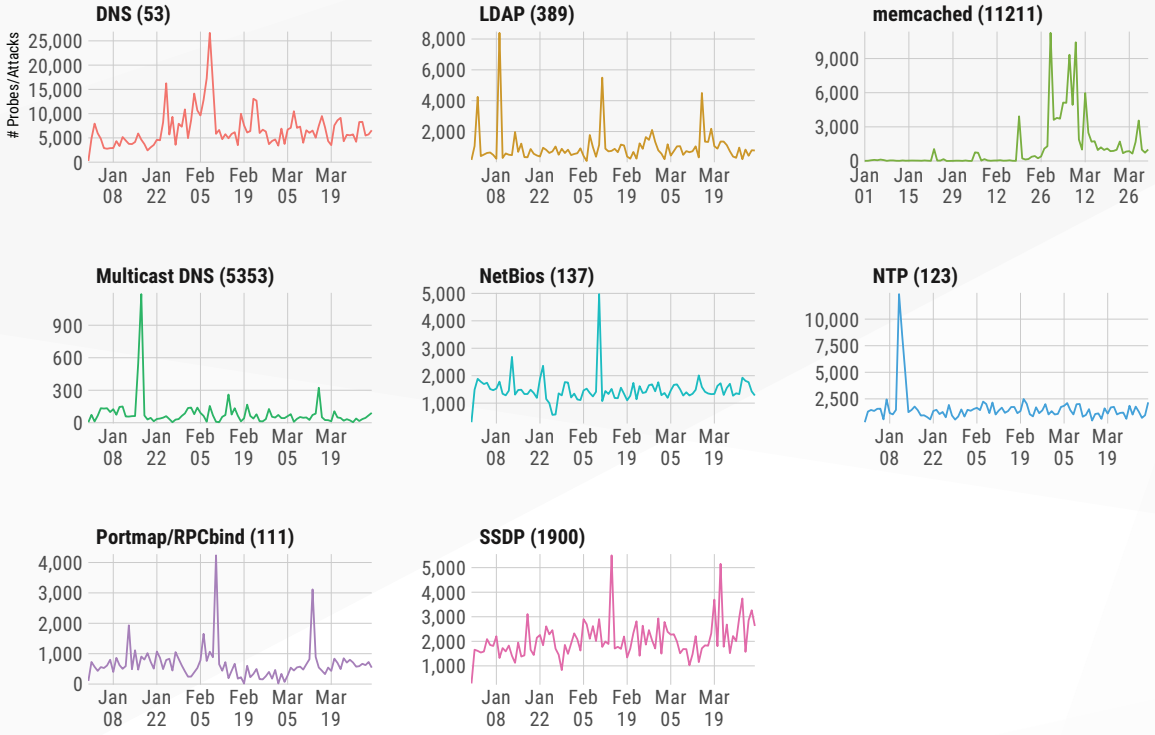

NOTE Free Y-axis Scale
Source: Rapid7 Project Heisenberg

Rapid7 Labs began detecting what can only be described as "inventory scans" across eight amplification DDoS ports starting in January this year (Figure 9). Each inventory scan was conducted (mostly) on a single day, and attackers used a mix of relatively new nodes and "the usual suspects" to perform these scans. Figure 10 shows these two distinct patterns for the top 20 IPv4-attributed countries for all of Q1 2018.

Figure 10: Part-time vs Full-time: Two Distinct Types of Attack/Probe Nodes

When Rapid7 Labs looked at the amplification DDoS inventory and attack data, two distinct patterns emerged: relatively new bot-nodes (many are IoT devices) that were used just to perform the inventory scans, and the 'usual suspects' that are regularly used in a diverse array of probes and attacks.

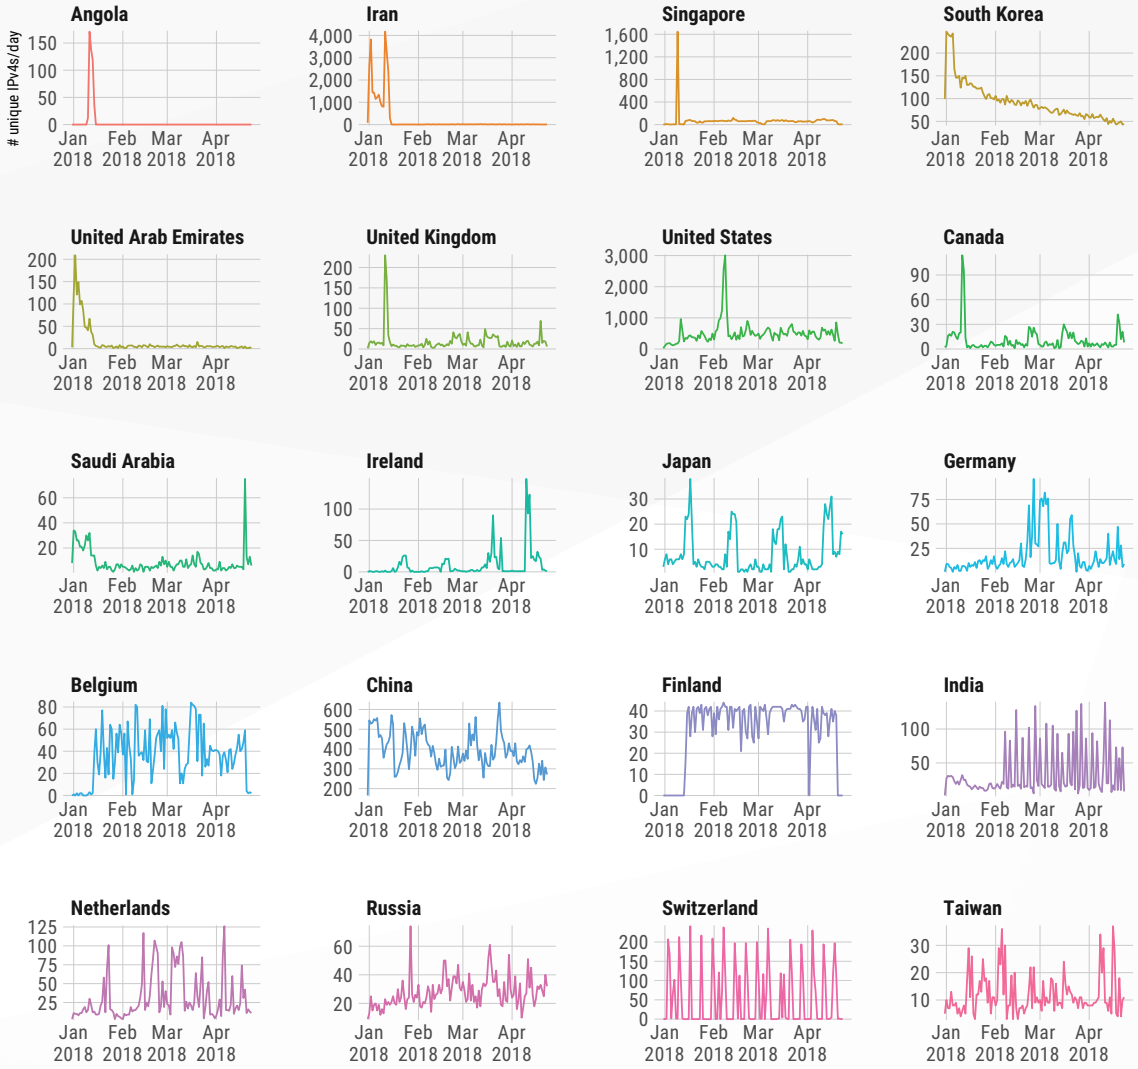

Broken down by source country; NOTE Free Y-axis Scale
Source: Rapid7 Project Heisenberg

DDoS attacks aren't just used to "punish" an organization. They can also be used to distract security, network, and operations teams while other malicious activities are being performed. It is vital that security teams practice for these "bad days" and distraction-oriented DDoS attacks to help ensure sufficient team capacity and organizational visibility.

Furthermore, given how easy it is to enlist a bot army to conduct these attacks, organizations that are concerned about application or network availability should make plans to invest in DDoS mitigation technologies or services and perform regular business continuity tests using DDoS attack scenarios to avoid some potential costly downtime. This is especially important in the wake of WebStresser.org's shutdown, as reported by Brian Krebs. Since they were a major player in the booter/stresser industry, this takedown is likely to cause some chaos in the DDoS-for-hire space. In fact, this takedown (and the events leading up to it) may be part of the reason we've been seeing these new sources of inventory scanning since January.

SLOW AND STEADY CONTINUES THE SMB RACE

Previous incarnations of this threat report have noted the popularity of Microsoft's Server Message Block (SMB) as an opportunistic attack target on the internet. Rapid7 Labs is keeping tabs on SMB and continues to see a steady increase in probe/attack levels along with the "daily grind" pattern as seen in Figure 11.

Figure 11: SMB Probes and Attacks Are Still Tracking Upwards

While the pace has slowed, there is a steady rise in daily, unique sources taking inventory of open/vulnerable SMB nodes and/or trying to compromise them.

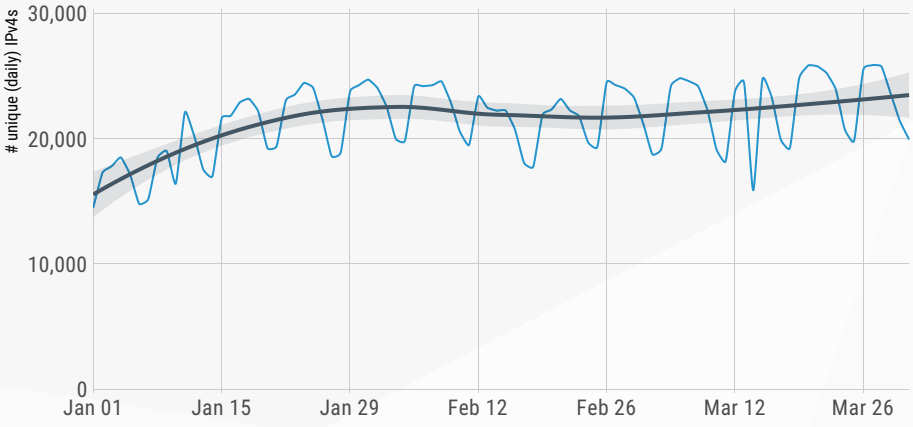

Source: Rapid7 Project Heisenberg

Attacks and probes are seen from virtually every IP geolocated country, but 12 of them make up the vast majority of the traffic and contribute heavily to this "day job" pattern as seen in Figure 12. The autonomous systems involved coincide with regions that have large deployments of compromised IoT devices or where ASN space is cheap and plentiful. The dip in Vietnam is not specific to one Vietnam ASN but is also not related to the February 27, 2018 undersea cable cut. Drop a note to research@rapid7.com if you have more information on what may have led to this temporary decrease in traffic.

Figure 12: Even the Bots Have Day-Jobs

A number of organizations suffered through DDoS-A attacks in Q1 2018.

Twelve IPv4-sourced countries account for the vast majority of SMB attack/probe traffic and most of that traffic has a clear weekly seasonal pattern.

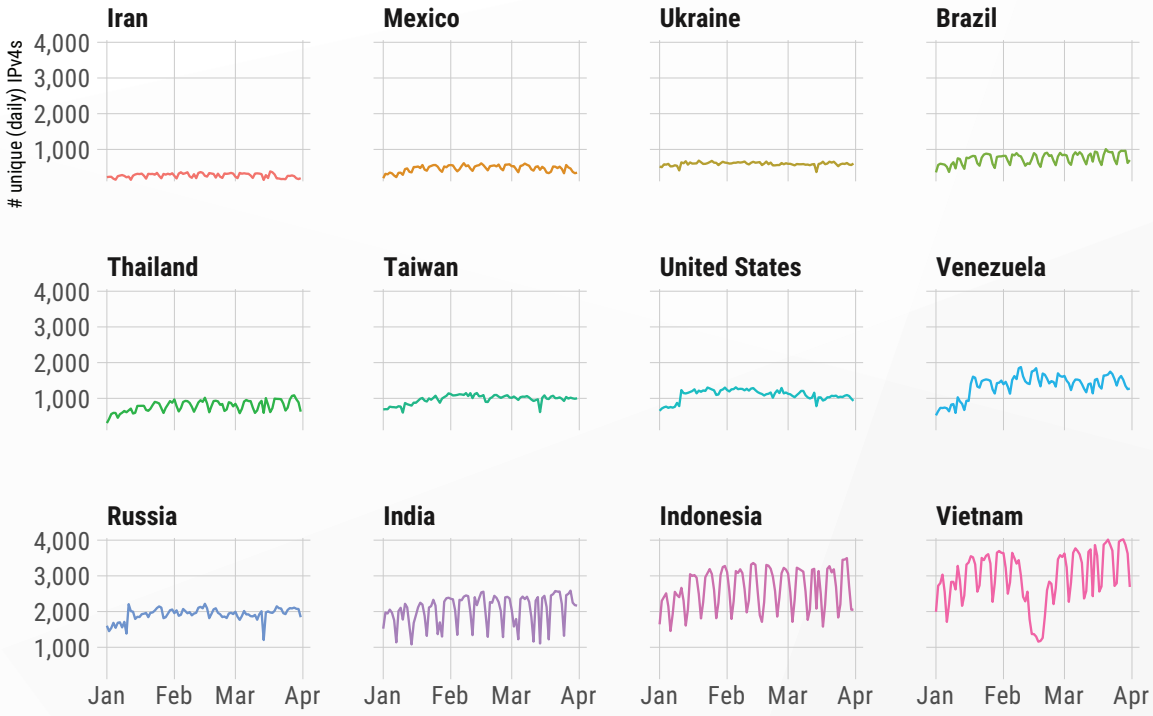

Source: Rapid7 Project Heisenberg

I SPY SMI

On March 28, 2018 Cisco released another **security advisory** for the Smart Install (SMI) feature of their Cisco IOS/IOS XE software. Rapid7 posted **an explainer** for the denial-of-service (DoS) and remote code execution (RCE) vulnerabilities documented in the advisory. This is not by any means SMI's debut performance, but we wanted to follow up in this report to help underscore the importance of not exposing SMI (port 4785) on the public internet. (In reality, organizations should only expose it on an internal, isolated network management subnet, if at all.) Figure 13 shows a spike in malicious activity immediately after the advisory was issued, and there has been a steady uptick in malicious traffic ever since.

Over 60% of the unique sources come from two autonomous systems: AS4837 (CNCGROUP China169 Backbone) and AS4134 (No.31,Jin-rong Street). Clearly, there is some serious Chinese interest—or adversary interest that is using Chinese IPv4 space for cover—in auditing the internet for vulnerable SMI installations. Organizations with a heightened sensitivity to attackers specifically from Chinese threat actors are well advised to double check their ingress firewall rules for port 4785 exposures.

Figure 13: Cisco SMI Malicious Probes/Attacks

Rapid7 Labs saw a significant uptick in the volume of malicious activity aimed squarely at Cisco SMI endpoints at the tail end of Q1 and well into Q2

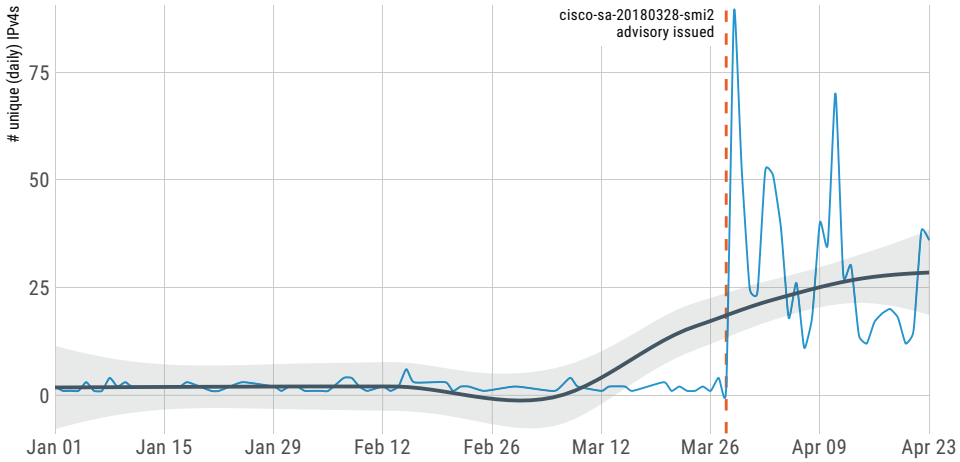

Source: Rapid7 Project Heisenberg + GreyNoise Intelligence

Organizations with a heightened sensitivity to attackers specifically from Chinese threat actors are well advised to double check their ingress firewall rules for port 4785 exposures.

YOUR 'SPRING CLEANING' TO DO LIST

The spring of 2018 certainly brings a mixed bag of challenges for information technology defenders. We've definitely seen some shifts in attacker behavior and identified new threat event patterns that may require organizations to dust out the corners of their defense plans.

First and foremost, keeping tabs on your organization's normal user behavior is critical in preventing credential-based breaches; our data from Q1 of 2018 shows that credential theft and usage—which in turn lead to suspicious logins—topped the charts for the most common threat events.

The memcached attack on GitHub was a harbinger of things to come for DDoS mitigation practices. Since each misconfigured memcached server commands such a massive bandwidth amplification factor, no organization on the internet is safe from even casual DDoSers until these vulnerable machines are taken down by their operators or their upstream internet service providers. The population of such amplifiers is shrinking, but there remain thousands of weaponizable servers active today. Double check that your organization isn't contributing to the arsenals of attackers.

SMB scans continue to hit high-water marks when defining the new normal levels of background attack scanning. These ever-increasing scanning levels are a testament to the effectiveness of EternalBlue as a tool in ransomware campaign execution. Organizations can help avoid being part of the problem by ensuring both inbound and outbound connection attempts to port 445 are blocked at least at the perimeter, and ideally, anywhere else where Windows networking isn't required to cross network segments.

Finally, the novel (and not so novel) attacks involving Cisco Smart Install (SMI) are attracting an unusual level of attention, as recorded by Rapid7's Heisenberg honeypot network. Again, this is a problem best solved by blocking any access to port 4785; even if your Smart Install implementations are fully patched and resistant to packet mangling tricks leading to remote code execution, this is not a service you want exposed to the internet at large.

That's a wrap on the first quarter of 2018. Stay tuned for the Q2 report in a few months.

APPENDIX A: METHODOLOGY

We gathered up closed and confirmed incidents from across a representative sample of our Managed Detection and Response (MDR) customers using our **InsightIDR** solution for the first quarter of 2018. Where possible, we've provided full incident counts or percentages; when more discrete information needed to be provided by industry we normalized the values by number of customers per industry. While we wanted to share as much information as possible, the precise number of organizations, industries, and organizations-per-industry is information no reputable vendor would publicly disclose.

As noted in situ, for this report we also incorporated data from both Project Sonar and Heisenberg Cloud. Raw Sonar scan data is available at **https://scans.io**, and you can contact **research@rapid7.com** for questions regarding Heisenberg Cloud honeypot data or any other findings or data used in this report.

The following table provides a full breakdown of the InsightIDR threat events and the threat event groups they belong in (as seen in Figure 6). Appendix B has the full, expanded listing of InsightIDR threat events.

IDR Threat Categories:

Dangerous User Behavior

- Account Visits Suspicious Link
- Password Set To Never Expire
- Network Access For Threat

Threat Probing

- Asset Connects To Network Honeypot
- Watched Impersonation

Threat Movement

- Account Authenticated To Critical Asset
- Lateral Movement Domain Credentials
- Lateral Movement Local Credentials
- Suspicious Authentication

Remote Entry

- Wireless Multiple Country Authentications
- Multiple Country Authentications
- Ingress From Non Expiring Account
- Ingress From ServiceAccount
- Service Account Authenticated From New Source
- Account Authenticated To Critical Asset From New Source
- New Local User Primary Asset
- Ingress From Disabled Account

Failed Access Attempt

- Authentication Attempt From Disabled Account
- Brute Force Against Domain Account
- Brute Force Against Local Account
- Brute Force From Unknown Source

Malicious Behavior On Asset Level

- Remote File Execution
- Log Deletion Local Account
- Harvested Credentials
- Log Deletion
- Virus Alert
- Network Access For Threat

Suspicious Behavior On Asset Level

- Malicious Hash On Asset

Malicious Behavior Network Level

- Advanced Malware Alert
- Protocol Poison
- Administrator Impersonation

Account Adjustment

- Account Privilege Escalated
- Account Enabled
- Account Password Reset
- Account Locked
- DomainAdmin Added

APPENDIX B: INSIGHTIDR THREAT EVENTS

EVENT	DESCRIPTION
Account Authenticated To Critical Asset	A new user authenticates to a restricted asset.
Account Authenticated To Critical Asset From New Source	A permitted user authenticates to a restricted asset from a new source asset.
Account Authenticates With New Asset	A permitted user is authenticating to an application from a new source asset.
Account Created	An account was created on a flagged asset.
Account Enabled	A previously disabled user account is re-enabled by an administrator.
Account Leak	A user's credentials may have been leaked to the public domain.
Account Password Reset	A user resets the password for an account.
Account Privilege Escalated	An administrator assigns higher level of privileges to the account.
Account Received Suspicious Link	A user receives an email containing a link flagged by the community or threat feeds.
Account Visits Suspicious Link	A user accesses a link URL identified as a threat from the Threats section or from other intel sources.
Advanced Malware Alert	An advanced malware system generates an alert.
Asset Connects To Network Honeypot	There was an attempt to connect to a network honeypot.
Authentication Attempt From Disabled Account	A disabled user attempts to access an asset.
Brute Force Against Domain Account	A domain account has failed to authenticate to the same asset excessively.
Brute Force Against Local Account	A local account has failed to authenticate to the same asset excessively.
Brute Force From Unknown Source	An unknown source has failed to authenticate to the same asset excessively.
Domain Admin Added	A user has been added to a privileged LDAP group.
First Ingress Authentication From Country	A user logs onto the network for the first time from a different country.
First Time Admin Action	An administrator action was used for the first time in this domain.
Harvested Credentials	Multiple accounts are attempting to authenticate to a single, unusual location.
Ingress From Disabled Account	A disabled user logs onto the network or a monitored cloud service.
Ingress From Non Expiring Account	An account with a password that never expires accesses the network from an external location.
Ingress From Service Account	A service account accesses the network from an external location.
Lateral Movement Domain Credentials	A domain account attempts to access several new assets in a short period of time.

EVENT	DESCRIPTION
Lateral Movement Local Credentials	A local account attempts to access several assets in a short period of time.
Log Deletion	A user deletes event logs on an asset.
Log Deletion Local Account	A local account deletes event logs on an asset.
Malicious Hash On Asset	A flagged process hash starts running on an asset for the first time.
Multiple Country Authentications	A user accesses the network from several different countries within a short period of time.
Multiple Organization Authentications	A user accesses the network from multiple external organizations too quickly.
Network Access For Threat	A user accesses a domain or IP address tagged in the Threats section.
New Local User Primary Asset	A new local user account was added to the primary asset of a domain user.
New Mobile Device	A user accesses the network from a new mobile device.
Password Set To Never Expire	A password of an account has been set to never expire.
Protocol Poison	Poisoning of a network protocol, such as via Responder, is detected.
Remote File Execution	Remote file execution has been detected.
Service Account Authenticated From New Source	A service account authenticates from a new source asset.
Spoofed Domain Visited	A user makes a DNS query to a newly registered internet domain.
Suspicious Authentication	A suspicious authentication was detected.
Virus Alert	A virus alert was triggered from an asset.
Watched Impersonation	A user authenticates to a watched user's account.
Wireless Multiple Country Authentications	A user logs onto the network using a mobile device from too many countries in a short period of time.

ABOUT RAPID7

Rapid7 powers the practice of SecOps by delivering shared visibility, analytics, and automation that unites security, IT, and DevOps teams. The Rapid7 Insight platform empowers these teams to jointly manage and reduce risk, detect and contain attackers, and analyze and optimize operations. Rapid7 technology, services, and research drive vulnerability management, application security, incident detection and response, and log management for organizations around the globe. To learn more about Rapid7 or get involved in our threat research, visit **www.rapid7.com**.

QUESTIONS?

Email us at research@rapid7.com