RAPID7

RAPID7 THREAT REPORT 2017 Q1

By Rebekah Brown, Threat Intelligence Lead, Rapid7, Inc.

Bob Rudis, Chief Security Data Scientist, Rapid7, Inc.

April 19, 2017

CONTENTS

INTRODUCTION

We had three goals in mind when developing the first edition of the Rapid7 Threat Report. First and foremost, we wanted to provide as clear a picture as possible of the threat landscape organizations faced during the first quarter of 2017. To this end, we've included composite and industry-level views of threat events across many industries.

While there's inherent value in just that threat portrait, we also wanted to show what a "day in the life" of a typical incident responder might look like, so we have created additional views by day, hour, and event type by industry to give you a glimpse into both the workload variety and volume facing these unsung heroes of cybersecurity.

Finally, as we examined the events of the past quarter, we've highlighted key takeaways that are applicable across organizations of every shape, size, and locale. We hope you find the report to be an informative and useful companion as you continue to develop your own detection and response programs.

We have created additional views by day, hour, and event type by industry to give you a glimpse into both the workload variety and volume facing these unsung heroes of cybersecurity.

THE 2017 Q1 THREAT LANDSCAPE

This report covers a representative sample of assessments from the first quarter of 2017 of both the generalized threat landscape, as well as unique threats that are more focused on certain organizations or industries. See the Methodology section in Appendix A for more information on how we collected and analyzed data used in the sections that follow.

The "threat landscape" is a moving, shifting form that will look different to different organizations—it all depends on where you are standing. Some people may be staring at a wide open grassland where the landscape is understood and the threats are easy to identify (though no less deadly), and others may be facing a dense jungle of hidden threats. While it is possible to develop a universal threat landscape that may be useful for general security research and developing high-level incident response frameworks, it is more useful for defenders to have a solid understanding of their own unique threat landscape.

Threat landscapes are defined by several factors, including: the type of data or access an organization has, the assets it has, the history of targeted threats, the sector and industry, and the impact that different threats would have on the organization. Figure 1 shows a view of the threat landscape by incident type and industry to demonstrate the diversity of threat actions across industries.

Figure 1: Incidents by Industry

Normalized across industries

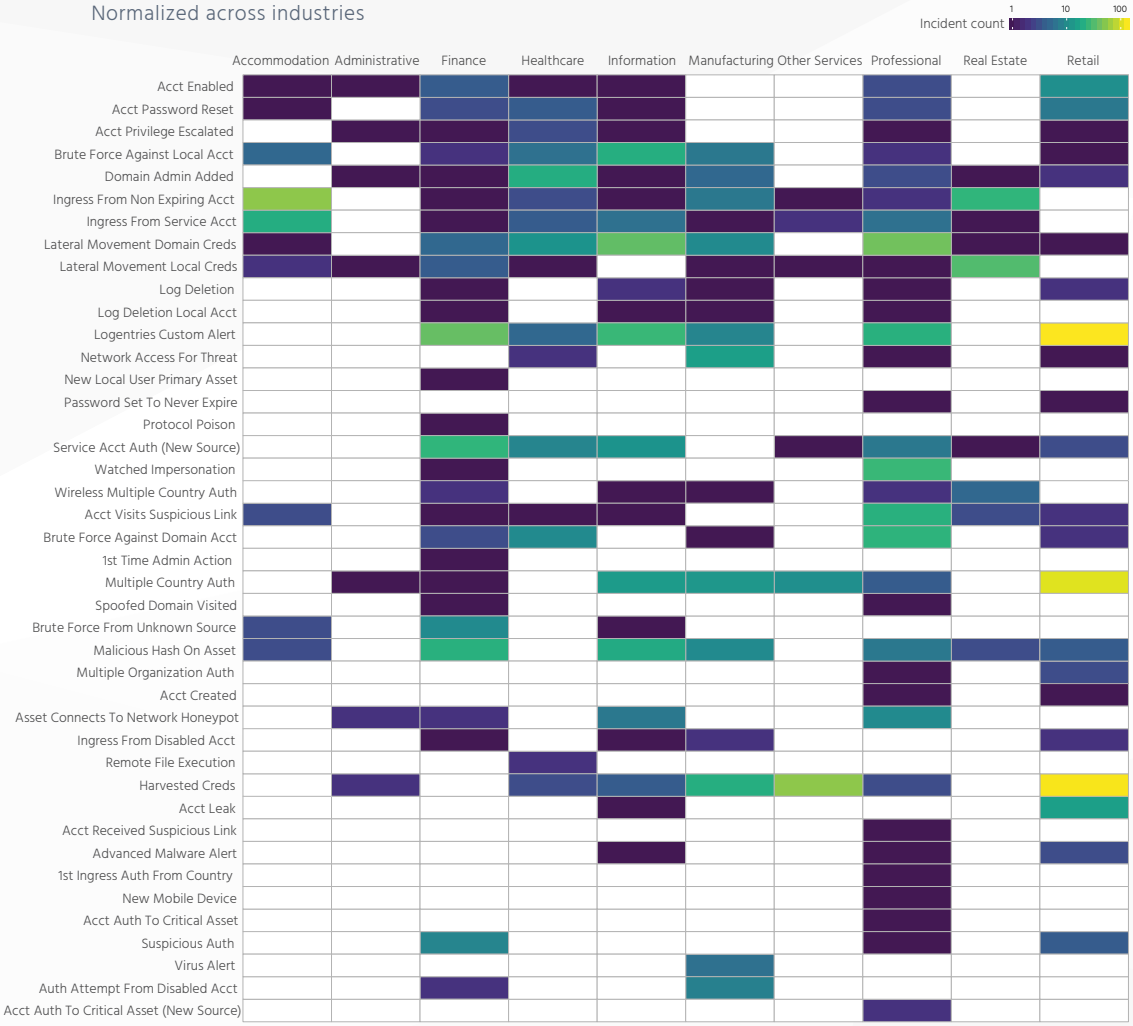

Source: Rapid7 Managed Detection & Response Incidents

Due to the variety of customers our managed detection and response (MDR) services and incident response (IR) services teams serve across multiple sectors – which face different threats and different priorities – our analysts need to understand threats across the board. They also need to understand the type of threats that are unique to sectors in a certain industry or with a certain type of data. Our analysts also need to assist each customer in taking the right action, so when we provide details on threats, we also need to provide details on how to respond to those threats. Knowing the enemy is a step in the right direction, but knowing how to deter and detect the enemy is even better.

Our focus is always on how to detect malicious activity, understand the threats, and respond and remediate when needed.

LESSONS LEARNED

Lesson #1: More is less. Less is more.

Alert fatigue is the name of the game for SOC analysts, and a great deal of emphasis is placed on reducing the amount of noise caused by false positives in alerting. There are several ways to deal with a high volume of alerts or a high number of false positives, and most of them involve a combination of tuning to ensure that the right things are alerting, and automation to handle true positive alerts that an analyst would resolve in the same manner over and over.

We saw a lower number of alerts that required a human analyst in 2017Q1 than we have traditionally expected. So, naturally, we patted ourselves on the back for not blowing up the analysts with false positives, before then attempting to see what else we could learn from the alerting trends.

Figure 2: Median incidents per day

Across all customers and industries

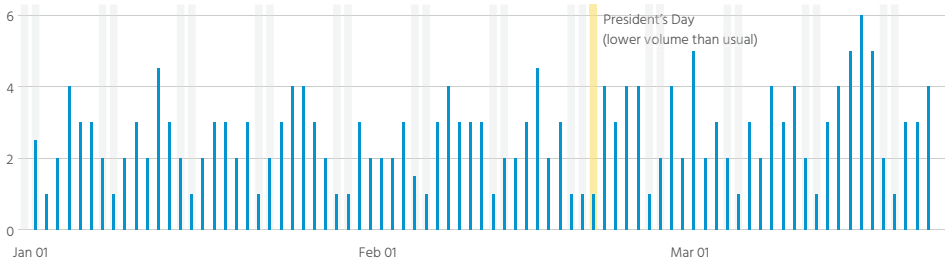

Source: Rapid7 Managed Detection & Response Incidents

Figure 2 shows a general pattern of fewer alerts on the weekends and an increased number of alerts during the middle of the week. Alerts stayed consistently low during the Presidents' Day (Monday) holiday as well. This may be because during the weekend or holidays there are fewer users to interact with attackers' lures, whether they come in the form of phishing emails, drive-by downloads, or other social engineering efforts. In fact, one of the primary lessons we pulled out from our data is that, across the board, attackers still rely on user interaction to carry out their attacks. This means that security measures a user can circumvent, wittingly or unwittingly, are not going to be completely successful at deterring attacks.

One quick note: This does not mean that attackers are not active during the weekend or breaks. It just means that those attacks that rely on user interactions to be successful will not actually generate an alert until the user takes some sort of action. A phishing email may have been sent on a Saturday, but if a user doesn't open it and download an infected PDF until Monday, then it will not alert until Monday.

Figure 3: Max incidents per day

Across all customers and industries

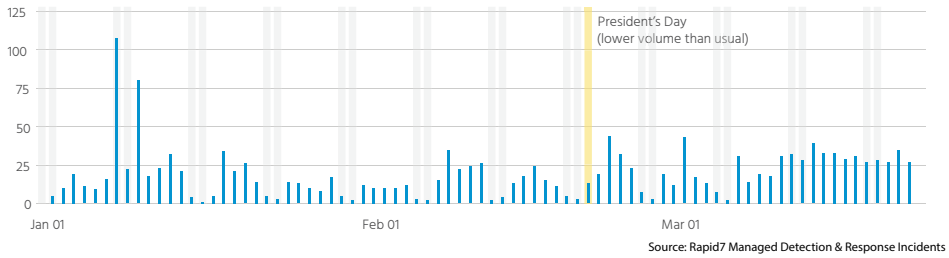

Source: Rapid7 Managed Detection & Response Incidents

We saw a lower number of alerts that required a human analyst in 2017Q1 than we have traditionally expected.

Unfortunately, having an overall steady median alert rate across multiple customers does not guarantee that there will not still be bad days in the SOC. When we look at the maximum number of alerts generated per day (Figure 3) we can clearly see some outlier days[1] where analysts had to scramble to respond to and understand a high number of alerts. Some of those were the result of a new detection method being brought online, and some were the result of a high number of opportunistic attacks. When spikes like these appear, it is important to understand why there is a deviation from the normal and what response was taken so that you can learn from these experiences, especially if it indicates a change in the threat landscape.

Figure 4: Hourly incidents heatmap per industry

Normalized across industries • Hours in GMT

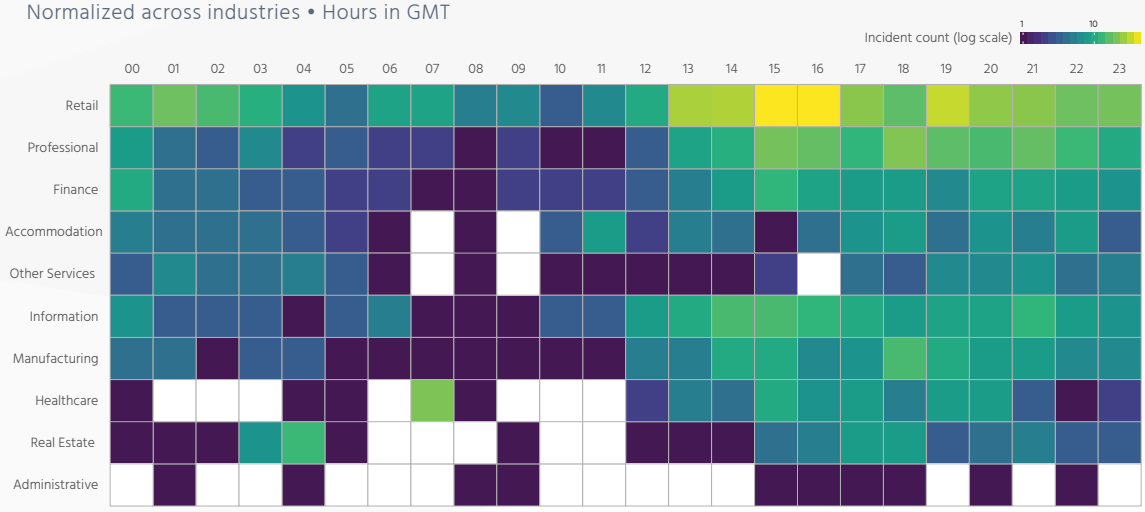

Source: Rapid7 Managed Detection & Response Incidents

The reliance on users may explain another trend we noticed in Figure 4, which is that attacks increase midday and later, which may correlate with the time most users are active. Analysts who are looking to quickly identify and respond to attacks should be aware of the hours of the day when the highest volume of alerts are generated. These high-activity hours will vary by organization, which is another reason why it is important to understand the unique aspects of the threat landscape for your organization.

Lesson #2: You find what you are looking for.

When we looked at what events trigger alerts (Figure 5) we found that most alerts come from known, bad activity, whether that is a particular behavior such as multiple concurrent logins from across the world, or malware or other artifacts left behind by attackers. The quality of these alerts depends on the inputs that are provided. Inputs based on low-fidelity indicators will generate far more noise than signal. If you base alerts on these inputs, guess what you will find when you investigate them: more noise than signal. When the input you are providing is something that you know is malicious, and you know that whenever it's seen that it will likely be malicious, then you will be able to generate alerts that can be trusted and acted upon.

In addition to these high-fidelity known malicious activities, a large number of alerts came from custom alerts, such as alerts based off newly identified activity, or activity that has been identified as relevant to only some organizations and is not applied universally. Not surprisingly, these custom rules generate alerts that are specific to the organization they were built for. Custom alerts require additional context because it is important for analysts to know why they were created and what they should do in response to them.

[1] While we called out Presidents' Day earlier, some of you might be curious as to what was happening on January 7, 2017. Take a peek at Figure 5 to see the cause as identified by InsightIDR.

Figure 5: Daily heatmap per incident type

Fill color shown as a percentage of incidents that day/column

Source: Rapid7 Managed Detection & Response Incidents

Not all custom alerts are good alerts, however. When there are high volumes of alerts generated from a single rule in a short time period, it either means that you have a big problem with your network, or you have a big problem with the custom alert. More often than not the problem is with the alert, but thankfully that is easier to fix than a big problem on your network.

Lesson #3: APT is dead, long live APT.

"Advanced Persistent Threats" are threats that are typically associated with state-sponsored activities. While there has been a lot of hype centered around APT-centric threats, it turns out that the majority of organizations were not impacted by APTs (APT is dead) … unless you happen to be in one of the industries that aligns with nation-state interests (government, manufacturing, aerospace, we are looking at you), in which case APT activity is alive and kicking (long live APT). This is one of the reasons why it is so important to understand your own particular threat profile and what the threat landscape looks like for you. APT activity—which is much more targeted and rare—results in a longer response period and requires a great deal more post-incident response in order to overcome the "persistent" part of the attack.

When we look at the majority of notable activity that we saw during 2017Q1, it was primarily identified quickly and remediated quickly. We did have a few instances where the same organization was impacted by the same threat multiple times before the root cause was identified, but in these cases it was often an untargeted attack that was successful multiple times rather than an attacker who learned from—and adapted to—a defender's actions. In the cases where we were dealing with an advanced persistent actor, the time to detect and the time to remediate increased significantly.

Understanding your organization's threat profile can help determine whether or not APTs are something that need to be accounted for in the threat landscape. If it is a consideration, then your incident detection and response plan should account for that, and detections should be specifically developed around those threats (see Lesson #2). If it is not part of the threat landscape that you need to worry about, then it is important to keep in mind the situations that may cause your threat profile to change, and in the meantime focus on the things that are impacting your organization.

Lesson #4: I feel the need, the need to Strut with speed.

There was a time when a 30-day patching cycle was sufficient to deter all but the most sophisticated attackers. That time, however, has passed. When a new vulnerability that is relatively easy for an attacker to exploit is identified, we may have days—not months—before systems are being compromised in a massive and often untargeted manner.

Figure 6: Apache Struts Vulnerability (CVE-2017-5638) Timeline

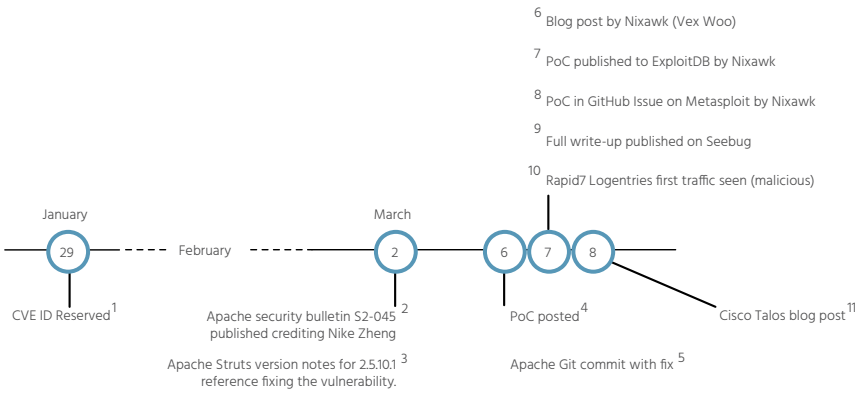

This is exactly what we saw in Figure 6 with the Apache Struts vulnerability (CVE-2017-5638). Our research first detected scanning and exploitation attempts against this vulnerability on March 7, and by March 9 we saw a massive increase in exploitation attempts. Weeks later the scanning and exploitation continued, and unpatched systems were significantly impacted. In some cases there were as many as nine distinct attackers attempting to exploit a single vulnerable system, using a variety of scanning methods, exploits, and post-exploitation tactics.

How do we function in a world where attackers move faster than we can remediate? There is not a one-size-fits-all approach, but understanding the threat landscape will go a long way toward shortening that time to respond and remediate critical issues. Attackers will not ALWAYS be faster. Not every vulnerability that is identified is cause for panic. Knowing when you need to act fast and when to stick with a slow and steady (and reliable) remediation plan is the key. The fact that the Apache Struts vulnerability was seeing widespread and indiscriminate exploitation just days after it was publicly disclosed is a good sign that swift actions should be taken. Understanding the threat, and being able to convey the risk that threat presents, can help get buy-in to expedite the patching process. If it is absolutely, physically impossible to expedite patching on systems then this would be another good reason to apply temporary, custom alerts that will identify attacks against the system (Again, see Lesson #2. Man, that one was a good one!). Alerts like these are likely to catch other, unrelated activity, but in a pinch, they will protect you while you work on your remediation plan.

When a new vulnerability that is relatively easy for an attacker to exploit is identified, we may have days—not months—before systems are being compromised in a massive and often untargeted manner.

CONTINUING EDUCATION

While we found no new, scintillating tales of high-stakes espionage campaigns in the cases we investigated during 2017Q1, each individual incident provided valuable information that facilitated increased understanding of the threat landscape and how it applies to different organizations and industries. Each measurement helps us identify the subtle changes in attacker methodologies and behaviors enabling the development of better ways to detect and respond to their actions.

As evidenced by this report and Rapid7's approach to research in general, we are strong believers in the value of transparency, information sharing, and mutual education in the ongoing endeavor to reduce opportunities for attackers. To this end, we will continue to share our threat intelligence findings for the benefit of the community; look for our report on this quarter's activity, out in early Q3. You can find information on our other research reports and findings at www.rapid7.com/research or on the Rapid7 Community blog.

APPENDIX A: METHODOLOGY

For this cardinal threat report, we gathered up closed and confirmed incidents from across a representative sample of our Managed Detection and Response (MDR) customers using our InsightIDR platform for the first quarter of 2017 (specifically, 2017-01-01 through 2017-03-22). 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.

Speaking of industries, we used the 2017 North American Industry Classification System (NAICS) (https://www.census.gov/eos/www/naics/) to classify the organizations. If you're unsure what NAICS sector your organization belongs to, you can usually find it with a quick internet search using "Your Org Name NAICS". If your industry is not represented here, you can contact us privately at research@rapid7.com as technical or other constraints may have prohibited including a specific industry view in our corpus this time.

The incident event types used in the visualizations come from our InsightIDR platform. The following is a list of the ones identified in this report with their descriptions:

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.

EVENT	DESCRIPTION
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.
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

With Rapid7, technology professionals gain the clarity, command, and confidence to safely drive innovation and protect against risk. We make it simple to collect operational data across systems, eliminating blind spots and unlocking the information required to securely develop, operate, and manage today's sophisticated applications and services. Our analytics and science transform your data into key insights so you can quickly predict, deter, detect, and remediate attacks and obstacles to productivity. Armed with Rapid7, technology professionals finally gain the insights needed to safely move their business forward. To learn more about Rapid7, visit www.rapid7.com.

QUESTIONS

Reach us at research@rapid7.com