FORTINET®

QUARTERLY

THREAT LANDSCAPE REPORT

Q2 2019

Table of Contents

Q2 2019 **Introduction and Overview**

Welcome back to our quarterly romp through the wild and crazy cyber-threat landscape. Q2 exhibited many themes and trends we've seen before, but we encountered plenty of new and noteworthy developments as we poured over intelligence collections. Here's a brief summary of what's on the menu for this quarter:

The Fortinet Threat Landscape Index

This barometer of threat activity across the internet hit its highest point ever in Q2 to close 4% higher than this time last year.

RobbinHood and Its (Un)Merry Men

The ransomware renaissance flourished in Q2 with attacks on Baltimore and other municipalities. Proceeds were not given to the poor.

RDP and the "BlueKeep" Blues

A spate of RDP vulnerabilities, including the infamous BlueKeep, reminds us that remote access services open a door for criminals too.

Upping the Ante on Anti-analysis

We examine a spam campaign that used novel anti-analysis and evasion techniques and discuss why this trend is one worth following.

Exploiting the Digital Supply Chain

Third-party risk is nothing new, but recent incidents exemplify the scope of exposures tied to a growing web of interdependencies.

Probing Smart Homes and Businesses

Between consumer IoT and ICS is a growing line of smart devices for home and small business use that has threat actors salivating.

Playbook Preview: Zegost

An infostealer active since 2011 has been upgraded with a plethora of capability upgrades. Our analysis will help you avoid being its next victim.

Exploratory Analysis: Vulnerability Research

28 zero days, vulnerabilities exploited in the wild, and time to exploitation for new signatures ... what more could you want?

The findings in this report represent the collective intelligence of FortiGuard Labs, drawn from a vast array of network sensors collecting billions of threat events observed in live production environments around the world. According to independent research[1], Fortinet has the largest security device footprint in the industry. This unique vantage point offers excellent views of the cyber-threat landscape from multiple perspectives, and we look forward to sharing highlights from that analysis with you in the pages that follow.

Threat Landscape Index

Q2 2019 Open: 1013 | Close: 1037 | High: 1037 | Low: 1004

The Fortinet Threat Landscape Index was developed to provide an ongoing barometer for aggregate malicious activity across the internet. Generally speaking, the TLI is based on the premise that the cyber landscape gets more threatening as more of our sensors detect a wider variety of threats at a higher volume. If the opposite is true, this indicates things are getting better. Perhaps most importantly, it shows the rate of those changes over time and helps draw attention to the forces behind them.

The TLI crossed a milestone this quarter, completing its first full year and giving us a good reason to review where we've been. Overall, the TLI is up nearly 4% from its original opening position at 1000. The high point during that year-long timeframe is 1037, which also happens to be the peak and closing point of Q2. Nothing like ending on a high note, huh?


Figure 1: Fortinet Threat Landscape Index (top) and subindices for botnets, exploits, and malware (bottom).

The upsurge at the end of Q2 was driven mainly by increased malware and exploit activity. Let's take a deeper look into a few specific detections behind that activity.

Chart Toppers

Figure 2 lists the top 10 detections for malware and IPS detections across Fortinet sensors in Q2. We'll start with malware since that index jumped the most at the end of the quarter. On that front, the most common by far were malicious files exploiting the CVE-2017-11882 memory corruption vulnerability affecting Microsoft Office documents.

Top 10 Malware detections

CVE_2017_11882	188k
Framer.INF!tr	116k
Agent.OAY!tr	63k
Abnormal.C!exploit	41k
ProxyChanger.ES!tr	38k
Agent.MUV!tr.dldr	37k
Agent.NIK!tr.dldr	34k
Heuri.D!tr	26k
Phish.EMW!tr	20k
RBot.BMV!tr.bdr	20k

Top 10 IPS detections

|	
ThinkPHP.Controller	3.3m
ThinkPHP.Request	2.5m
PHP.Diescan	2.4m
Apache.Struts	1.9m
Joomla!.Core	1.9m
MS.IIS	1.2m
Drupal.Core	1.2m
HTTP.URI	1.2m
MS.Windows	900k
HTTP.Header	874k

Figure 2: Most common malware variants and IPS detections by device in Q2 2019.

The 2017 date on this vulnerability is a bit misleading, as that marks the date it was publicly disclosed. The flaw's existence traces back 17 years prior to its discovery, meaning it remained unknown for quite some time. Conspiracy theories suggest it was known and used surreptitiously by various government agencies around the world during that long period of dormancy, but that's never been proven. We ran recent samples through a sandbox and confirmed there was nothing new or special about the Word/RTF files that execute upon being opened. It's been popular among various actors and campaigns over the last couple of years, and this latest activity looks to be nothing more than the latest chapter in that saga. Figure 3 shows the global setting for that story.


Figure 3: Global detections of CVE-2017-11882 malware samples in Q2 2019.

One of these stories in the Q1 2019 TLR focused on exploits targeting content management systems. ThinkPHP, Joomla, and Drupal are major players in that space, and Figure 2 shows they continue to attract attention among cyber criminals. If you use them, don't neglect them!

Not shown in Figure 2 are two exploit detections against the Open Dreambox and Spree Commerce platforms that topped our list of "major movers" over the quarter. These echo a theme we address in the *Exploiting the Digital Supply Chain* story below. These platforms and plugins appear to be all the rage among criminals looking to exploit third-party dependencies throughout the online transaction chain.

Featured Q2 Updates

We pored over our intelligence collections, sensor data, and media headlines from Q2 2019 with an eye toward identifying a handful of interesting topics and trends to share. There's no formal or common criteria for inclusion other than they all feature a different view of the threat landscape our analysts found noteworthy. We hope they assist you in seeing over the Q3 horizon and beyond.

RobbinHood and Its (Un)Merry Men

Multiple high-profile incidents last quarter served to highlight the rising impact of ransomware attacks for organizations that are not prepared to deal with them. In May, an attack on the city of Baltimore disrupted critical services for weeks and forced officials to implement manual workarounds for handling real estate transactions, utility payments, property taxes, and other critical functions. Baltimore officials, acting on the advice of the FBI, refused to pay the approximately $100,000 the attackers wanted as ransom and ended up spending more than $18 million on recovery efforts.

Fortinet's analysis of RobbinHood, the ransomware used in the Baltimore incident, showed it as designed to attack an organization's network infrastructure and likely distributed via weaponized remote desktop applications. The malware's capabilities include the ability to disable Windows services that prevent data encryption and to disconnect from shared drives.

Several other municipalities and government entities experienced similar attacks in Q2 2019. One of them was Riviera Beach, Florida, which spent three weeks trying to recover data that was encrypted in a ransomware attack, before paying $600,000 to the attackers for the decryption key. Another victim, Lake City, Florida, paid $490,000 to cyber criminals to avoid disruption following a similar attack. Many view such payments as likely to only encourage more ransomware attacks in the near term, even though victims without proper data backup and recovery processes might see it as the only viable option.


Figure 4: Devices detecting Ryuk variants.

Some reports identify the ransomware in the Florida attacks as Ryuk, which surfaced last year. It uses several evasion tactics, including destroying its encryption key and deleting shadow copies on an infected system. Researchers attribute Ryuk primarily to targeted attacks, and the low number of detections in Figure 4 supports that position. It's likely distributed via spear phishing or brute-force attacks on RDP services.

Last quarter's ransomware attacks continue a trend that began last year. Cyber criminals are moving away from mass-volume, opportunistic ransomware attacks and are increasingly focusing their efforts on organizations that they perceive as having both the ability and the incentive to pay big ransoms. It's common in these targeted attacks that cyber criminals have gained access to the victims' networks and conducted considerable reconnaissance before deploying their ransomware on carefully selected systems. Some report a sharp increase both in the average ransom payments that organizations are making to get their data back as well as in disruption-related costs.


Figure 5: Downward trend in ransomware detections (number of devices) during Q2 2019.

During the second quarter we also observed a new ransomware sample called Sodinokibi (aka Sodin) surface that could soon become a major threat for enterprise organizations this year. Functionally, Sodinokibi is not very different from a majority of ransomware tools in the wild. What makes it troublesome is the fact that it exploits a recently announced critical vulnerability in Oracle's WebLogic Server (CVE-2019-2725) that allows for arbitrary remote code execution. The impact of this could be severe because a vulnerable system can get infected without the victim doing anything to trigger it. Organizations that have not yet patched Oracle WebLogic Server versions 10.3.6.0 and 12.1.3.0 should do so immediately.

Takeaway: The attacks on Baltimore, multiple cities/local governments in Florida, and elsewhere in Q2 serve as a reminder that ransomware, while declining in overall volume, continues to pose a serious threat for organizations going forward.

RDP and the "BlueKeep" Blues

Few security vulnerabilities received as much attention or caused as much concern last quarter as "BlueKeep," a critical flaw in the Remote Desktop Services function in multiple older versions of Windows. The flaw (CVE-2019-0708) allows an unauthenticated user to connect to a vulnerable system via Microsoft's proprietary Remote Desktop Protocol (RDP) and take control of it to steal credentials and data and for planting ransomware and other malware.

Microsoft and others, including the U.S. Department of Homeland Security, have repeatedly warned organizations about the severity of the flaw and advised prompt patching. Even so, at the end of Q2 2019, internet scans showed there were more than 800,000 systems with RDP services exposed to the internet that were unpatched and vulnerable to attacks. Research Fortinet conducted in June unearthed several vulnerable systems on Microsoft Azure data-center IP ranges. When contacted, Microsoft advised us the IPs likely belonged to third-party customers and not to the company itself. Our research leads us to believe that other cloud service providers and their customers are likely impacted in a similar fashion.

Attackers have leveraged RDP quite heavily in recent years to access remote Windows systems and execute a variety of malicious actions on them. However, Microsoft has described BlueKeep as being especially of concern since it is "wormable" and allows malware to spread autonomously from one vulnerable system to another in the same manner as the notorious WannaCry ransomware in 2017. Microsoft disclosed BlueKeep in May and released patches for it including for Windows versions that it no longer actively supports.


Figure 6: Proportion of devices detecting RDP exploit attempts in Q2 2019, including scans related to BlueKeep (MS. Windows.RDP.Channel.MS_T120.Remote.Code.Infection and MS.Windows.RDP.CVE-2019-0708.Remote.Code.Execution).

Since then, there have been multiple reports of proof-of-concept code and exploits becoming privately available for BlueKeep. The DHS announced in June that it had successfully tested a remote code execution exploit for BlueKeep against a legacy Windows system. A Metasploit module has purportedly been developed but is being held back from release until after the first outbreak. Microsoft and security experts believe it is only a matter of time.

According to the FBI, there has been a steady increase in attacks exploiting remote admin tools since mid-to-late 2016 largely because of the growth in underground markets selling access to RDP credentials. Often the attacks have been facilitated by organizations using weak passwords to protect access to RDP, by running outdated versions of the service, or by providing open access to the default RDP port (TCP 3389). Recent examples of malware that cyber criminals have distributed via badly secured RDP services—and that we have reported on—include Dharma (aka CrySiS), SamSam, and GandCrab ransomware strains.

Takeaway: BlueKeep is a reminder for organizations to secure RDP services. Best practices for mitigating risk include disabling RDP on systems that don't require it, using strong passwords and account lockout to protect against brute-force attacks on RDP, applying available patches and updates to address known vulnerabilities, and enabling network-level authentication.

Upping the Ante on Anti-analysis

Many modern malware tools incorporate features for evading antivirus and other threat-detection measures. Common examples of these anti-analysis techniques include routines that enable the malware to detect when it is running within a sandbox environment or an emulator; functions for disabling security tools on an infected system; and the use of junk data to make disassembly harder. MITRE lists more than 60 anti-analysis and evasion techniques—some new and some old—that attackers can employ to slip past an organization's defenses.

A macro that we observed being used in a major spam campaign in Japan last quarter is a good example of how adversaries are constantly using and tweaking these anti-analysis techniques to stay ahead of defenders. The spam campaign involved the use of a phishing email with an attachment that, in this case, turned out to be a weaponized Excel document with a malicious macro. Our analysis showed the macro had attributes for disabling security tools, executing commands arbitrarily, causing memory problems, and ensuring that it would only run on Japanese systems.

Like many other malicious software, the rogue macro in the Japanese spam campaign was designed to look for certain Excel-specific variables at multiple points during execution to ensure it was running within an Office Excel environment and not in an emulator. One Excel property that it looked for in particular—xlDate variable—was something that we haven't observed before in other malware. Interestingly, the variable appears to be undocumented in Microsoft's documentation—or at least we were not able to find it.

```
15    End Function
16
17  ⊟Private Sub Workbook_Open()
18        'debug.print
19  └     If xlXmlExportValidationFailed > 0 Then oceran
20    End Sub
21
22  ⊟Function lowsharts()
23  └     lowsharts = timefortime(3 - 2, 2 - 1)
24    End Function
25
26  ⊟Function Betal()
27        fdsgfadsrff436tgdfzf33546s = 1 + (Application.International(LittlePeace) - 1)
28  └     Betal = fdsgfadsrff436tgdfzf33546s
29    End Function
30
31  ⊟Function LittlePeace()
32  └     LittlePeace = ((xlDate))
33    End Function
34
```

Visual length : 1,506   lines : 61   Ln : 19   Col : 35   Sel : 27 | 1   Windows (CR LF)   UTF-8   INS

Figure 7: Example of anti-analysis used by the macro. Not too many macros will actually check for Excel-specific variables such as xlXmlExportValidationFailed. By doing this, the authors have ensured that the macro is executed within an Office Excel environment. This means that macro emulators may fail if they do not properly emulate specific Excel variables.

Anti-analysis techniques like these aren't new but their usage appears to be growing. In June, security researchers spotted a new variant of the Dridex banking Trojan that successfully evaded several traditional antivirus tools by using 64-bit DLLs with file names of legitimate Windows executables. The file names and associated hashes changed each time the victim logged in, making it hard for signature-based antivirus tools to spot the malware on infected host systems. June's Dridex variant also took advantage of a known weakness in the Windows Management Instrumentation Command-line (WMIC) utility to bypass application whitelisting measures and execute malicious VBS code embedded within an XSL file.

Last quarter we encountered several reports of downloaders with sophisticated defense-evasion techniques built into them. One example is AndroMut, a downloader that the Russian-speaking TA505 group used last quarter in a campaign targeting individuals working at financial companies in the U.S. and elsewhere. AndroMut's anti-analysis features included sandboxing and emulator verification and checks for mouse movement and debuggers. At least two other downloaders—Brushaloader and a new version of JasperLoader—were reported in Q2 2019 as having similarly advanced evasion mechanisms including location verification capabilities and sleep timers for delayed execution.

Takeaway: The growing use of anti-analysis and broader evasion tactics pose a challenge to enterprise organizations and underscore the need for multilayered defenses that go beyond traditional signature and behavior-based threat detection.

Exploiting the Digital Supply Chain

Multiple major incidents in the second quarter highlighted the growing threat to organizations from attacks via supply chain partners and other trusted third parties.

Figure 8: Weekly detections of Cryxos variants associated with a Magecart campaign in Q2 2019.

One was an operation we reported in May involving more than 185,000 payment cards stolen from numerous ecommerce sites using a lightweight JavaScript card skimmer detected as JS/Cryxos. PWS!tr. We observed steady activity for this variant early in Q2 and then a large spike in June (see Figure 8). Our analysis showed that in the early stages of the campaign, the cyber criminals were able to nab nearly 40,000 payment cards per month. Most victims were U.S.-based, with other pockets in Australia, Britain, France, and Italy.

Magecart is an umbrella term for multiple criminal groups that have been stealing credit and debit cards from high-traffic web businesses by embedding card-skimming software on their sites. Often the JavaScript skimmers have been inserted into third-party components that the sites rely on for functions as varied as content management, visitor tracking, customer support, and payment services. Magecart made headlines in 2018 for stealing card data from hundreds of websites globally including those belonging to major companies such as British Airways, Ticketmaster, and OXO. British Airways was recently fined nearly $230 million for the breach. The fine is the biggest to be levied under the EU's General Data Protection Regulation to date and underscores the high costs associated with third-party risks.

In June we uncovered another campaign using similar tactics to steal card data from ecommerce sites. In this case, our investigation began when we found a JavaScript skimmer, called Inter, disguised as a traffic tracker for a website. We found Inter to be designed to intercept and capture credit card details that users entered into payment forms on compromised websites. Our analysis showed Inter to be highly configurable to fit each buyer's needs and available in underground markets for $1,300 per license.

Digital card skimmers in third-party components are certainly not the only threat. In Q2 2019 we also encountered several reports where attackers gained access to an organization's network by first breaking into a system belonging to a supplier or other third party with trusted access to the network.

China's APT10 group reportedly broke into systems belonging to at least eight very large IT service providers and used the compromised systems as launch pads for attacks on the customers of those vendors. APT10 is believed responsible for hacking the networks of scores of managed service providers around the world for the same reason. In another example, cyber criminals compromised a major hardware vendor's automatic software update server and used the system to distribute malware instead. The authors of the "ShadowHammer" campaign were specifically targeting systems belonging to a small, carefully selected subset of the hardware vendor's customer base. But thousands of other computers suffered collateral infections as well.

Takeaway: Such reports highlight the need for effectively managing risk from third parties. Under mandates such as PCI DSS and HIPAA, organizations using these third-party services are primarily responsible for protecting the data. A failure to do this can have serious consequences.

Probing Smart Homes and Businesses

Our last featurette this quarter offers a perspective that we haven't covered extensively in prior reports. We've written a lot about threats in the Internet of Things (IoT) and the constant probing of consumer devices like home routers, IP cameras, printers, etc. On the other side of the spectrum, we've discussed threats against critical infrastructure and analyzed exploits targeting industrial control systems (ICS) and supervisory control and data acquisition (SCADA) technologies in great detail.

But somewhere between your home printer and critical infrastructure is a growing line of control systems for residential and small business use. These smart systems garner comparably less attention than their industrial counterparts. But that may be changing.

Our data muse for this story comes from increased scanning activity we observed targeting Schneider Electric devices. Since Schneider Electric is one of the leading manufacturers of industrial control devices, an initial assumption is to categorize such activity as probing for access to operational technology (OT). And there's ample historical precedent for that conclusion. The threat group dubbed Xenotime has recently been carrying out broad scans of dozens of U.S. power grid targets. Xenotime is believed to be the group behind the Triton attacks on Schneider's Triconex SIS controllers at a Middle Eastern oil company a couple of years ago. And earlier this year, they reportedly hit about half a dozen North American oil and gas targets.


Figure 9: Proportion of devices detecting probes targeting Schneider U.motion devices.

This activity, however, doesn't target Schneider's Triconex controllers. The detections that caught our attention involved the company's U.motion line, which they describe as a building management solution. The related signature triggered in 1% of organizations, which may not seem like much, but is much higher than we typically see for Schneider's (and other manufacturers') ICS or SCADA products. Figure 8 shows daily prevalence immediately jumped more than tenfold after the signature was deployed in mid-May and then another tenfold a month later.

Once attackers identify hosts/victims that have a U.motion panel (normally an open webpage), the exploit in question is easy to perform via a simple PHP webpage post method. Successful exploitation grants control of the Schneider device, and more importantly, access to any number of other devices under management—environmental controls, security cameras, safety systems, etc. It doesn't take much to imagine the harm a resourceful criminal could do with that kind of access. And it doesn't take much to understand why the security of smart residential and small business systems deserves elevated attention.

Takeaway: The fact that the U.motion exploit went viral so quickly demonstrates criminals are watching closely for opportunities to commandeer control devices in homes and businesses. Unfortunately, cybersecurity in these venues is usually overlooked and underfunded, especially outside the scope of traditional IT systems.

Playbook Preview: Zegost

Zegost, also known to operate under the alias of Zusy or Kris, is an infostealer originating in China that has been active since 2011. Since its inception, Zegost has added a plethora of updates. Offensive capabilities were greatly improved by acquiring the capability to use specific PowerShell actions that download the infostealer the moment a victim's mouse moves over a specific piece of text. The added ability for Zegost to clear its own event logs gave the infostealer long-term evasion capabilities, granting more time to move laterally within the victim's network. Zegost also gained the ability to access and record the victim's webcam. A previous update went so far as to enable it to utilize COM programming, an uncommon feature in malware.

Like other infostealers, the main objective of Zegost is to gather information about the victim's device and exfiltrate it. It will hunt for OS versions, analyze the speed and quantity of processors in the victim's machine, check for an internet connection, and look for the RDP port number. Zegost will also hunt for a login number for QQ, which is a Chinese chat client, a feature that corresponds with its Chinese origin.

Compared to other infostealers, Zegost is uniquely configured to stay under the radar, making it far more of a long-term threat compared to its contemporaries. The malware accomplishes this by clearing its own event logs as well as evading runtime conflicts by creating a mutex, which it checks to ensure only a single version of itself is running. Another recent development in Zegost's evasion capabilities was a command that kept the infostealer "in stasis" until after February 14, 2019, after which it began its infection routine. Another way Zegost stands out from other infostealers is the ability to launch processes via a window that can be hidden in order to avoid detection. Zegost has also been known in the past to exploit Adobe Reader vulnerabilities, specifically CVE-2013-0640, which allows for remote arbitrary code execution.

Zegost has become the favored infostealer of many threat actors and has been used in a variety of campaigns in the past. The Linux botnet Mr. Black used Zegost to gain access to routers in past campaigns and assimilate them into the botnet for the purpose of DDoS attacks. The infostealer was previously unleashed against government organizations in Nepal as well as multiple Vietnamese targets. The most infamous use of Zegost occurred in 2015 when the Italian offensive security company, The Hacking Team, was compromised by the infostealer and had a list of exploits used by the team leaked to the general public.

Currently, Zegost is the cornerstone of a spear-phishing campaign against a Chinese governmental entity providing statistical analysis on a number of subjects. The motives for this campaign are unclear at this time. While Zegost hosting infrastructure is based mainly in China, third-level domains for the infostealer have been observed outside of the country. Access the full Zegost playbook via our online viewer.

Exploratory Analysis: Vulnerability Research

As the title implies, this report examines the numerous types of threats that dot the cyber landscape. But since that landscape is heavily impacted by vulnerabilities in hardware and software (and even wetware), we devote a lot of attention to that topic as well. In fact, vulnerability research represents a major strategic focus for FortiGuard Labs. Not only do these efforts make our own products more secure but they also benefit the broader community in several ways. This section explores some of our vulnerability research activities and accolades from Q2 2019.

Zero-Day Research

Our experts examine many third-party products and software applications daily for weaknesses and exploitable vulnerabilities. When found, the FortiGuard Labs team notifies the software or product vendor of the vulnerability and creates protective measures that can be delivered to our customers.

Table 1 lists the products for which we disclosed vulnerabilities in Q2 2019. The team also discovered another 33 zero days that do not yet have a vendor fix available, and so full details have not been published.

Product	Count	Product	Count
Adobe Magento	1	LiveZilla Server	7
Adobe Shockwave Player	7	Microsoft Office	1
Cisco WebEx	3	Microsoft Windows	5
Ignite	1	Oracle	1
Keysight EMPro	1	RocketChat	1

Table 1: Zero-day vulnerabilities disclosed in Q2 2019.

As a testimony to these efforts, Fortinet was recognized by Microsoft at BlueHat 2019 in Shanghai for being among the top five vulnerability reporting organizations. Learn more about these vulnerabilities and our other zero-day research efforts at https://fortiguard.com/zeroday.

Predicting Exploitation

Not only does Fortinet conduct vulnerability research but we also support others doing it as well. In June, researchers from Virginia Tech, the Cyentia Institute, and RAND presented a paper at the Workshop on the Economics of Information Security (WEIS). The paper develops a machine-learning model for predicting vulnerability exploitation that performs demonstrably better than existing prioritization methodologies like the Common Vulnerability Scoring System (CVSS).

Prior related research based prediction models on vulnerabilities with proof-of-concept or exploit code that is publicly available, rather than those that have actually been exploited in the wild. Aware of our vast array of devices monitoring exploit activity across the internet, the researchers asked Fortinet to provide sanitized data on exploits in the wild. We were happy to oblige. Due to this expanded dataset provided by Fortinet, the researchers had nearly 3x the amount of exploited vulnerabilities for use in training and testing prediction models than any prior research! Based on that data, the researchers determined that 5.5% of all 100,000+ vulnerabilities contained in the National Vulnerability Database have been exploited in the wild. Access the full paper here.

Clocking Exploitation

In addition to understanding the likelihood of exploitation, we're also interested in studying the timelines of vulnerabilities exploited in the wild. As vulnerabilities and exploits become known (by us or others), we create detection signatures and deploy them to FortiGuard devices. As you can imagine, this is a constant, race-against-the-clock process.

How much time passes between the release of a signature and the first (or peak) detection of exploit activity? To study that, we took a sample of signatures released in Q2 and examined exploit activity for the first 30 days after release. Figure 10 records the results.

Figure 10: Prevalence of exploitation during the first 30 days of signature deployment.

Figure 10 contains numerous interesting observations, and we'll highlight a few of them here. First, it's plain why signature development and deployment must be done quickly. Many exploits are active in the wild by the time the signature is deployed, which is a major reason we do not release vulnerability information until the vendor creates a fix. No need to give attackers a head start. It's also evident that the path to peak exploitation varies substantially. Some blaze immediately. Others take a while to get to full intensity. Still others never go beyond a slow burn or flicker intermittently.

Understanding exploitation timelines like this helps nullify the first mover advantage that attackers tend to have over defenders. And honestly, that's a major goal for this entire report. Better information on the threat landscape around us informs better strategies to protect the interests and assets of our organizations. Thanks for exploring it with us, and we look forward to heading back out into the wilds again with you next quarter.

References

[1] IDC Worldwide Security Appliance Tracker, March 2019 (based on annual unit shipments).

FURTINET®