INTERNET SECURITY REPORT

Q3 2023

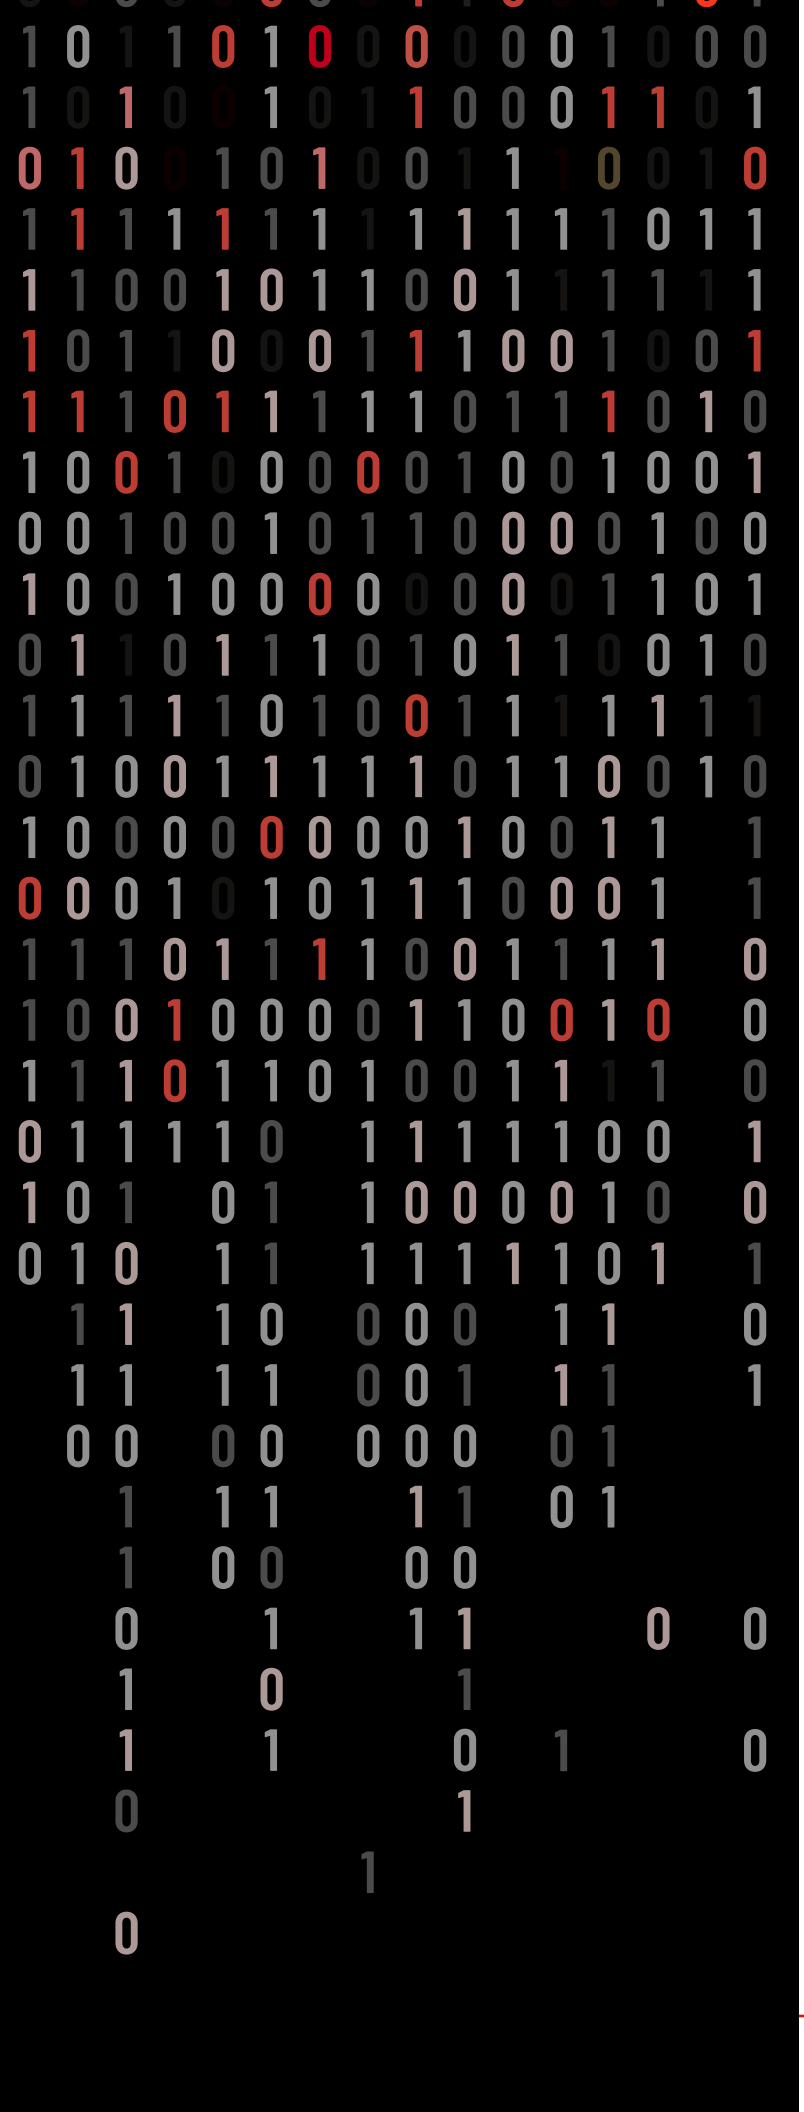

CONTENTS

The Firebox Feed™ provides quantifiable data and trends about hackers' latest attacks, and understanding these trends can help us improve our defenses.

INTRODUCTION

"True wisdom is acquired by the person who sees patterns, and comes to understand those patterns in their connection to other patterns - and from these interconnected patterns, learns the code of life..."

~Hendrith Vanlon Smith Jr, The Wealth Reference Guide: An American Classic

Put simply, we do this report to find patterns in the current cyber threat landscape so that you can figure out the best defenses against those common attack patterns.

Sounds so basic, but if you think about it, analyzing patterns is the core of science. The stories about great scientific minds discovering new things often start with them seeing some pattern in real life and wondering why they occur like Newton wondering why an apple always falls to the ground. Confirming a scientific hypothesis is an exercise in recording data to find patterns, making controlled changes to those systems, and watching to see if or how the pattern changes in result. In short, recognizing and analyzing patterns over time is one of the ways humans learn new things, which helps us figure out how to manage them.

This report is all about analyzing the cyber threat patterns happening on the Internet over time, so you can understand them and make the right defense choices. However, for something to become a pattern and not just a one-time occurrence, you need to see it happen repeatedly over time. Over the last two quarters, we have made significant changes to the methodologies we use to gather, curate, and report on our threat landscape data in our report. The downside to these changes is all our historical patterns and references no longer remained relevant to the new way we present our data. We could not compare the last two quarters' results to the patterns seen in previous quarters since different methodologies will create different results. That prevented us from making wider, quarter-over-quarter (QoQ) and year-over-year (YoY) discoveries and comparisons.

However, this quarter we can start paying attention to long-term patterns again. Now that we have three quarters under our belt with our new methodology, we have some historical data to compare. This allows us to get back to discovering long-standing threat patterns, making more QoQ and eventually YoY comparisons again, and recognizing when those patterns break. So expect future reports to return to longer-term pattern analysis.

Patterns and new methodologies aside, if you are new to this report, you may wonder how we're tracking and analyzing the threat landscape at all? Every quarter, we aggregate threat telemetry from tens of thousands of WatchGuard network appliances and millions of endpoint products whose owners have opted to share this data with us. We analyze this data to identify threat patterns, such as the most common or widespread malware, or the most prominent network attacks. We do so in hopes of giving you some ideas of what cybercriminals have been doing, and to find patterns that might suggest how those threat trends might evolve in the future. Of course, once we identify these trends and patterns, we also suggest the best defense strategies you should deploy to shield against them.

More specifically, our Q3 report includes:

EXECUTIVE SUMMARY

During Q3, network malware detections rose overall. Raw malware detection is up 14% and the sophisticated and evasive threats stopped by our behavioral detection service, APT Blocker, increased a whooping 129%. On the flip side, we saw malware detected over encrypted connections (TLS) drop considerably compared to the past two quarters. Some malware highlights include an email-based dropper family called Stacked taking multiple spots on our Top 10 malware and Top 5 encrypted malware lists, and a Chinese website delivering some commoditized adware and password stealers.

Unlike Q2, Network attacks increased significantly during Q3. Of note, ProxyLogin a critical Microsoft Exchange vulnerability that could lead to remote code execution topped our network attack list and accounted for 10% of all network attacks. You should have patched this critical flaw long ago, but if you haven't this should act as your wake-up call.

Compared to network-based malware detection, endpoint malware is down QoQ, at least for unique malware detections. This doesn't necessarily mean the raw volume of malware is down, only that the amount of unique malware variants we detected during Q3 is down. That said, the number of indicators of compromise (IoC) our endpoint products detected was up this quarter. From a delivery-vector perspective, script-based malware has dropped significantly from a few quarters ago, even though it remains number 1. However, Windows file-based malware delivery has risen and become a strong second. In any case, attackers leverage both malicious scripts and commonly exploited Windows files for their livingoff-the-land (LotL) methods. Whether one is up or the other is down, both suggest that LotL malware delivery is gaining popularity with threat actors.

That is just a taste of some of findings from the report. Below find some additional Q3 highlights:

- **Total network-based malware detections were up 14%** with malware detection from the APT Blocker service in particular up an amazing 129%, which suggests sophisticated and evasive malware continues to grow.
- **Our "per Firebox" malware results for various network malware detection services:**
  - **Average total malware detections per Firebox: 1,343** (14% increase)
  - **Average malware detections by GAV per Firebox: 507** (2% decrease)
  - **Average malware detections by IAV per Firebox: 474** (6% decrease)
  - **Average malware detections by APT Blocker per Firebox: 362** (129% increase)
- We extrapolate that if all the Fireboxes reporting to us had all malware detection services enabled, we would have had **100,925,107 malware detections during Q3 2023**. Note, that number only represents the Fireboxes that have opted into sharing data with us, it would be significantly higher if it included all active Fireboxes in the world.
- **Endpoint ransomware attacks increased nearly 90%.** On the surface, endpoint ransomware detections appeared down in Q3. Yet the Medusa ransomware variant, which emerged in the Top 10 malware threats for the first time, was detected with a generic signature from the Threat Lab's automated signature engine. When factoring in the Medusa detections, ransomware attacks rose 89% quarter over quarter.
- Surprisingly, **malware hiding behind encryption (TLS) dropped to 48% during Q3.** After many quarters of increase, this is a surprising new trend. Nonetheless, we still highly recommend scanning encrypted traffic since almost half the malware arrives over encrypted traffic.
- **Zero-day malware accounts for 69% of all malware.** As a reminder, we define zero-day malware as malware that evades signature-based protection, only detected by machine-learning malware models or behavioral analysis. This is a big increase compared to Q2. Furthermore, zero-day malware detected over TLS increased to 76%.
- **An email-based dropper, Stacked, comprised four of the Top 5 TLS malware detections in Q3.** All but one of the variants in the Encrypted Top 5 contained the dropper malware Stacked, which typically arrives in emails with malicious attachments that claim to include an invoice or other important document for review. Two of the Stacked variants also appeared in the Top 10 malware detections.
- **Network attacks increased 16 percent quarter over quarter (QoQ).**
- **ProxyLogin was the most attempted network attack during Q3.** As a reminder, this was a critical, remote code execution vulnerability against Microsoft Exchange servers that you should have patched long ago.
- **Our endpoint protection products blocked 171 unique malware variants per 100k machines.** This represents a steep 83% decline from Q2.

- **Script malware delivery continues to decline,** while other malware vectors, including Windows files, are up.
- **Threat actors increasingly leverage legitimate remote management tools in their attacks.** Among the new top phishing domains, we found a tech support scam that would result in a victim downloading a pre-configured, malicious version of TeamViewer, which would allow an attacker full remote access to their computer.

Those are only some of the findings from this quarter's report. If you're interested in more details about why we may have seen these patterns in quarterly changes, want to learn specifics about highlighted attacks and malware, and most of all, want the best defense tips against these trends, be sure to read on.

FIREBOX FEED STATS

WHAT IS THE FIREBOX FEED?

This section is built entirely on anonymized data that Firebox customers have opted in to sharing with us. The threat intelligence we receive allows us to view the specific malware and exploit activity that threat actors are using against small and midsize organizations worldwide.

In this section, we detail the high-level quarter-over-quarter trends while also diving into the specific top threats that generate either the most alert volume or impact the most unique networks. Through these lenses, we identify trends in the categories of malware or network attacks targeting WatchGuard customer networks and use that information to prescribe specific tips for a strong defense.

We break the Firebox Feed up into three main sections built off telemetry from five security services running on Firebox appliances:

Gateway AntiVirus (GAV): Signature-based malware prevention

IntelligentAV (IAV): Advanced AI-based malware prevention

APT Blocker: Sandboxed, behavioral-based malware prevention

Intrusion Prevention Service (IPS): Network-based client and server exploit prevention

DNSWatch: Domain-based threat prevention

HELP US IMPROVE

Our data comes from Fireboxes in our Firebox Feed and the more Firebox admins that provide the anonymous data the better we can make our reports. If you configure your Firebox to do so, we will have more accurate information in this report to apply to your network. So please configure your Firebox to enable device feedback by following these steps.

1. Upgrade to Fireware OS 11.8 or higher (we recommend 12.x)
2. Enable device feedback in your Firebox settings
3. Configure WatchGuard proxies and our security services, such as GAV, IPS, APT Blocker, and DNSWatch, if available

Average combined total malware hits per Firebox

1,343

Our average malware hits per Firebox, for devices that have all three services

Basic Gateway AntiVirus (GAV) service

507

Basic malware detections decreased slightly by **2%**

APT Blocker (APT)

362

Advanced evasive malware detections jumped by **129%** from the previous quarter

IntelligentAV (IAV)

474

IAV hits dropped by **6%**

GAV with TLS

109

TLS detection by GAV decreased **86%**

APT Blocker with TLS

585

Encrypted evasive malware dropped **31%**

TLS malware %

48%

48% of malware detections came from an encrypted connection

MALWARE TRENDS

The data we receive from Fireboxes opted into our Firebox feed allow us to audit the threat landscape based on what the devices block. The anti-malware service results specifically give us a great idea of what malware trends happen every day. Like any good audit, we see an overview of how users configure the security of their network. We know for example that only 20% of Fireboxes scan for malware within encrypted connections, yet 48% of malware arrives through encrypted connections. We hope these networks have additional layers of security, such as endpoint detection and response (EDR) solutions, to protect against malware arriving over the encrypted connections they are not monitoring. However, we know some networks don't. In this section we will discuss all the malware trends from the quarter, including more detail about the security risks from not scanning encrypted traffic.

We saw a return of three Linux-based malware families in Q3. Two of these revolve around the Linux.Lucifer botnet. We also saw an increase in malware detection in the Americas (AMER) almost matching the number of hits in Europe, the Middle East and Africa (EMEA). Much of this volume came from the #1 malware detected, Adware.Generic.3112968, which we primarily saw in AMER. Finally, we found two threat groups that pretend to provide legitimate software, but upon deeper inspection both provide hacking tools to their users.

We'll cover all this and more, but first let's look at the high-level data for the quarter.

We not only use the Firebox Feed data to build this report, but also to identify areas where we can improve our WatchGuard products' security. If you would like to help with these improvements, please enable **WatchGuard Device Feedback** on your device.

Top 10 Gateway AntiVirus (GAV) Malware Detections

We review our GAV results and combine the known results from IAV to create the Top 10 table. This table shows the detected malware families with the most hits overall. This year the Top 10 table makes up more than half of malware overall.

Similar to Q2, we saw several overlapping families of malware in our Top 10 table. We separate these variants to help better explain what we see and to provide an easier-to-understand view of our data. Each malware sample in a larger family will exploit the same vulnerability or at least perform similar actions. Each variant in the malware families will likely access the same domains and run the same commands.

Again, we saw three malware variants targeting Linux servers. And again, the Top 10 list included different variants of the malware Linux.Lucifer. This quarter, Linux.Generic.314124 replaced Linux.Generic.295484 as the dropper for the Coinminer Linux.Generic.13476. We saw an exponential increase in Linux.Lucifer with the dropper taking the second spot in the Top 10 chart. For more details on this malware family see our **Q2 report**.

During Q2, a completely new malware family, Lazy.360502, made our Top 10 list. It contains a trojan that will download the adware 2345explorer, which we also discussed in the **Q1 report** from this year. We'll describe it in more detail later in this section and uncover some interesting connections it has with the Vidar password stealer.

Threat Name	Malware Category	Count	Last Seen
Generic.3112968	Adware	1,360,963	Q2 2023
Linux.Generic.314124(Linux.Lucifer)	Dropper	987,992	new
GenericKD.68079600	Adware	629,365	new
Linux.XORDDoS.AT	Dropper	588,993	Q2 2023
Linux.Generic.13476(Linux.Lucifer)	Coinminer	322,318	Q2 2023
Lazy.360502	Dropper	112,830	new
Logan.581	Password Stealer	93,821	Q2 2023
Zusy.255797	Win Code Injection	92,715	Q2 2023
Stacked.1.12	Dropper	71,713	new
Stacked.1.7	Dropper	58,516	new

Figure 1. Top 10 Basic Malware Table

Top 5 Encrypted Malware Detections

Only about one in five Fireboxes scan encrypted traffic. So, when we look at the previous Top 10 malware table, we don't see the full extent of the encrypted malware. For the Top 5 Encrypted table we isolate the malware detected over encrypted connections to show what we think is the more common and complete picture of the malware landscape. If more of our customers took advantage of our free feature to decrypt HTTPS, we believe this list would better represent the top malware.

We saw a significant change in the Top 5 Encrypted malware this quarter. All but one of the variants contain the email malware Stacked. We also saw some overlap with the previous Top 10 table. Heur.LShot.1 malware attempts to gain access to your system to inject code. We found one sample would load ransomware after infecting the victim through RDP access.

Threat Name	Malware Category	Hits
Stacked.1.12	Dropper	71,712
Stacked.1.7	Dropper	58,516
Stacked.1.8	Dropper	30,248
Heur.LShot.1	Win Code Injection	15,614
Stacked.1.26	Dropper	15,246

Figure 2. Top 5 TLS Malware Table

Top 5 Widespread Malware Detections

Our widespread list represents the threats that the most amount of Fireboxes detect. Our previous lists were just the highest by volume, but those may only affect a smaller subset of Fireboxes, perhaps in different regions. For widespread threats, we are looking at the malware that is commonly detected on many Fireboxes.

Malware often downloads other malware, sometimes called child payloads, and in a scenario where you have good network perimeter defense, nonetheless an unprotected device becomes infected, the malware on that device will continue to try downloading additional malware repeatedly until someone notices. We do our best to remove excessive, anomalous detections from our report, but scenarios like these can skew results. To guard against this, we have removed repeated infection attempts from this table's findings and only highlight initial infections. We believe the Top 5 Widespread table most accurately identifies what malware threats most networks admins must protect against.

Once again, the most widespread malware family, Adware.JS.Agent.FM, targets the most Fireboxes. Further down, a dropper Trojan.JS.Agent.USF, targeted 60% of our users in India. Trojan.JS.Agent.USF links to malicious websites that will attempt to download more malware or phish the victim. We also see several AMER users targeted by that malware family. Finally, we see three different Office malware families mostly target EMEA. AMER and Asia Pacific (APAC) should also watch for Office exploits as well, since we still see a high number of targets in those regions too.

Top 5 Most-Widespread Malware	Top 3 Countries by %			EMEA %	APAC %	AMER %
Adware.JS.Agent.FM	Dominican Republic - 38%	Indonesia - 36.36%	New Zealand - 33.04%	17.78%	10.81%	24.70%
Exploit.MathType-Obfs.Gen	Greece - 26.81%	Germany - 23.84%	Hong Kong - 20.51%	16.51%	6.23%	5.15%
Exploit.RTF-ObfsObj-Dat.Gen	Greece - 27.03%	Spain - 26.03%	Germany - 22.81%	16.11%	6.61%	4.24%
Trojan.JS.Agent.USF	India - 59.86%	Australia - 16.85%	Thailand - 16.56%	6.92%	10.30%	10.93%
Trojan.Groooboor.Gen.37	Germany - 21.03%	Hong Kong - 18.59%	Poland - 13.76%	11.93%	5.59%	3.21%

Figure 3. Most-Widespread Malware Table

Geographic Threats by Region

We just mentioned how Trojan.JS.Agent.USF heavily targets India. Since malware targets each region differently, we like to try to shed some light on where we see the most malware. We then normalize this data with the amount of Fireboxes in a region, so sales trends don't skew the results.

Comparing the chart below to Q2, we saw very little difference in the regional distribution of malware in Q3. EMEA (32.32%) and AMER (32.29%) almost evenly split malware distribution with close to identical percentages per Fireboxes. APAC (35.39%) saw the most malware per Firebox, with 3% more malware than the other two regions. Comparing it to previous quarters, during Q1 and Q2 of this year AMER had the lowest share of detections. Between Q1 and Q2, EMEA and APAC switched between having the highest number of detections. We also saw an additional 166 detections per Firebox from Q2 to Q3. We attribute this change to an increase in detection in AMER, particularly the malware Generic.3112968.

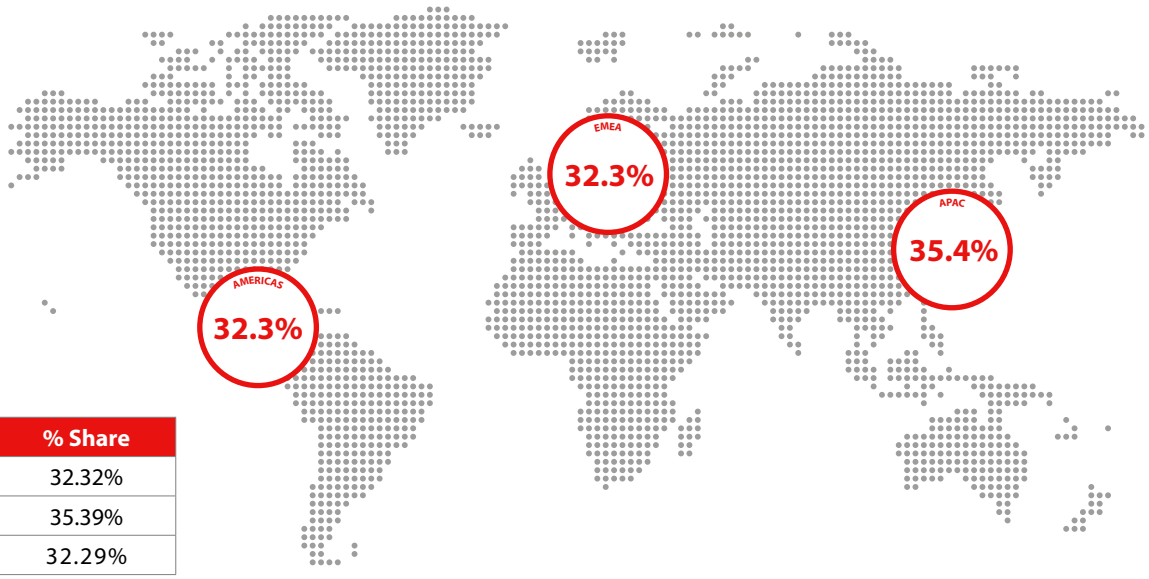

Region	% Share
EMEA	32.32%
APAC	35.39%
AMER	32.29%

Figure 4. Geographic Threats By Region

Catching Evasive Malware

Behavioral detection solutions, like APT Blocker, can often identify brand new zero-day malware based on the malicious things it does. Using advanced sandbox testing, it detonates the suspected malware to learn what the file does. If for example, APT Blocker finds the file connects to a malicious URI, that becomes one of the weighted indicators that the file is malicious. APT Blocker doesn't just rely on one behavior, but pays attention to hundreds of potentially malicious weighted indicators and adds the result to make a decision. Regular antivirus typically just looks for known patterns and doesn't always catch new stuff.

We saw an increase in zero-day malware in Q3 with both encrypted and non-encrypted traffic. 31% of all reported detections come from zero-day but when looking at just encrypted traffic 76% of malware detections are zero-day.

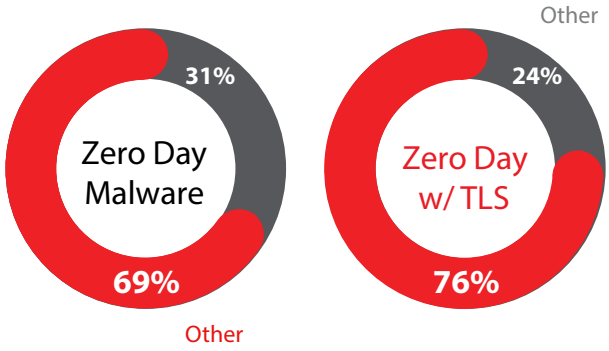

Figure 5. "Zero Day" Malware

Zero-day perimeter detection works well with host-based advance EPDR. While EPDR covers servers and workstations that have the service installed, you can't put EPDR on your printer. You also can't install EPDR on servers you don't know about. Layering these together provides the best protection for every host, server, printer, and IoT device in the office.

Individual Malware Sample Analysis

Gen:Variant.Lazy.360502
The dropper Lazy.360502 currently downloads the malware 2345Explorer we discussed in our Q1 report this year. We found that in the past, the server infected by Lazy.360502 downloads a variant of the Vidar password stealer. We previously suspected servers that provide the malware 2345Explorer will also spread other malware and this is the case here.

The specific IP that Lazy.360502 connects to also hosts the Chinese website upzxt[.]com. Upzxt.com main page currently provides a bootable hack tool to break into Windows systems and access filesystems.

Figure 6. Windows Hacktool

At one time, upzxt.com provided a credential stealer related to Vidar Pro Stealer. This malware service allows anyone to infect the victim's computer and steal credentials from them. Just like any other service, you pay as you go.

Figure 7. Vidar Pro Stealer Management Interface

Lazy.360502 and the 2345Exploer ecosystem will continue to spread adware, malware, and hack tools. We recommend avoiding sites with adware. If you see signs of adware on your device, you should run a complete scan for malware on the device as well.

Stacked 1.8
We saw several email-based malware variants, coming from the Stacked family, arrive over encrypted connections during Q3. One of the sample emails we analyzed contains an attachment called ORDER 20231015.rar. This file opens a Win-code-injector with the filename Lmguzn.exe and a generic icon. We found the Lmguzn.exe will capture user inputs, read the clipboard, dump OS credentials, hijack browser sessions, and perform other malicious activities related to stealing passwords.

If you don't recognize the sender or the invoice in an email doesn't open the email. Even if you do know the sender, we recommend checking with the sender, not through email, that they meant to send you an attachment before you consider opening it.

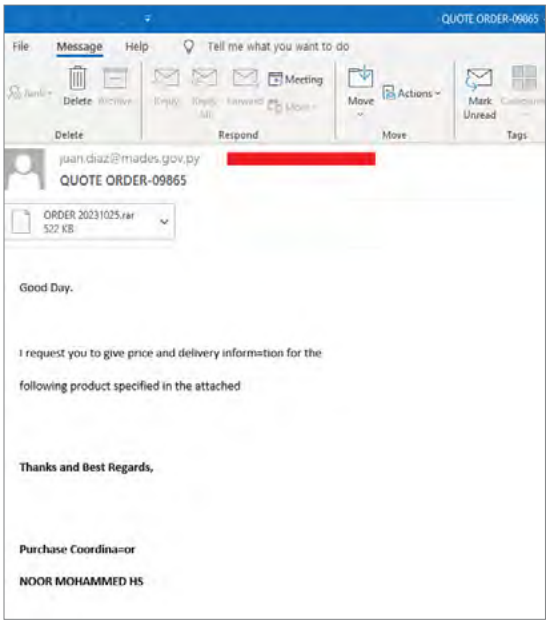

Figure 8. Stacked 1.8

Zum.Androm.
Looking further down the Firebox malware detections list, we found interesting results from Zum.Androm, which ended up being the remote access trojan (RAT) called Remcos. This trojan can access the victim's desktop, keystrokes, clipboard, and hijack browser sessions.

The not-so-subtle domain name https://breakingsecurity[.]net claims to provide penetration testing tools and the Remcos malware. It doesn't seem like they are marketing to the pen testing community though. We see Remcos marketed as an administrative tool. This mixed message raises a red flag.

Recent research shows that the group or at least one person in the group engaged in malicious activities. We will leave it up to the reader to determine the developers' intentions, but we assume the worst.

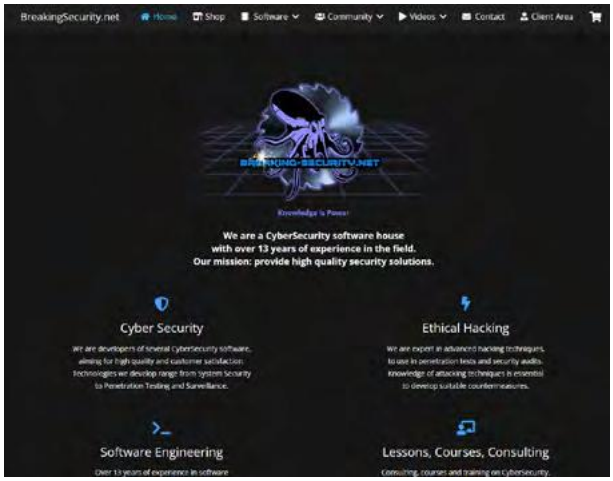

Figure 9. Breakingsecurity Homepage

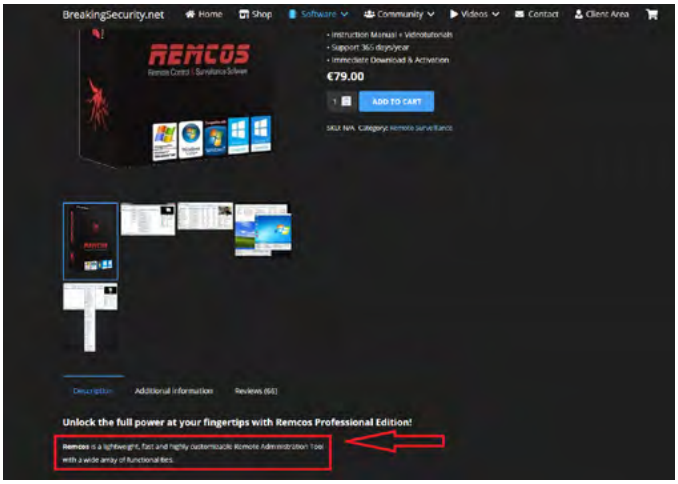

Figure 10. Remcos Admin Tool

Even if the developers only intended Remcos for pen-testing, seeing this tool in any network should raise red flags as it has been seen in multiple malware campaigns. We know of two threat groups that use this malware: the Gorgon Group, with connections in Pakistan, who target government organizations in the UK, Spain, Russia, and the US. APT33, with ties to Iran target aviation and energy sectors in the US, Saudi Arabia, and South Korea.

Network Malware Summary

Make sure the partners and groups you work with have good reputations in the security field. Groups that provide adware tend not to have high standards on who use these products and services. As we saw in this section, adware often comes with malware. We recommend not accessing websites that spread adware.

You could use hacking and penetration testing tools for testing your network security, but also for breaking into networks. Who it is that uses these tools makes all the difference. If you want to test your network security, only work with tools from groups or companies you trust and have a good reputation. We would avoid anyone who uses Remcos to test your network.

NETWORK ATTACK TRENDS

The WatchGuard Firebox Intrusion Prevention Service is a signature-based protection against known network attacks. This service is capable of identifying and blocking exploit attempts against old and new vulnerabilities. Many of the attempted attacks target vulnerabilities from the past 5-10 years, but some stretch further back. In this section we will review several aspects of the network attack data we receive from the Firebox IPS service. The main data we review are the Top 10 signatures by detection volume. Entries in this list are often tied to widely known exploits such as the ProxyLogon, the Microsoft Exchange Server vulnerability from 2021. In fact, for a second time, ProxyLogon sits at the top with the most detections per signature for the quarter, and this time with a larger percentage of detections among our 520,080 total detections this quarter. Additionally, attacks against this vulnerability were one of the most-widespread attack detections we had in the quarter. Total detections across all signatures increased by 16% when compared to the previous quarter. The total volume hovering around half a million may be considered a "new normal" after we updated our outlier data criteria a few reports ago to present more accurate information. If you look at Figure 11, you can see a relatively similar volume of detections between Q3 2022 and Q1 2023. Another new normal, tied to the total volume, is the overall average detections per Firebox. It went from 99 last quarter to 104 now. Further down in this section, we will break down the average detections per Firebox by region.

There were 389 unique signatures detected this quarter, a 6.49% decrease from the previous quarter. A quick glance at Figure 11 will make clear that a decrease of that size is in line with prior quarters. A statistic to hammer in this view is that since Q3 2020, the average combined increase is only 0.1%.

There are four new signatures among the Top 10 signatures by volume this quarter. Three of those four have never been present in the Top 50 signatures in the past few years – a quick rise, unlike other signatures that gradually made their way up into the Top 10. This can be seen in Figure 13 displaying the placement of the Top 10 signatures since Q3 2022. Signature 1138800 and 1059958 in 1st and 2nd place, respectively, are prime examples of the gradual rise.

The other six signatures in the Top 10 list are familiar. Five of them were present last quarter and one most recently in Q1 2023. An additional statistic we began tracking last quarter is signatures present among the Top 50 signatures by volume, that have never been on that list (going back several years). This doesn't mean they are newly published signatures, but that they have never generated a significant number of detections. Although it can be a new signature by publication as well. An important section along with the Top 10 signatures by volume is the most-widespread attacks. We track the top 5 and include the top three countries affected along with the regional percentages as well.

Drupalgeddon, Squid, ProxyNotShell, Quagga, and NOP Sled – the IT and security industry/community always finds a way to make technologies sound more fun than they are. Squid Proxy sounds better than something like "Reliable Proxy." All or some of those products and attack nicknames may be familiar to you already. Each of them is associated with one of our signatures this quarter, which you'll read about in this section.

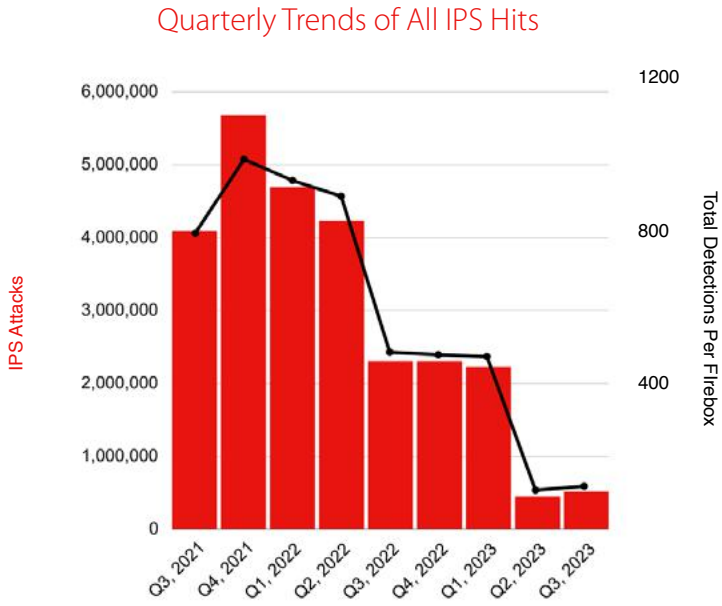

Figure 11. Quarterly Trends of IPS Hits

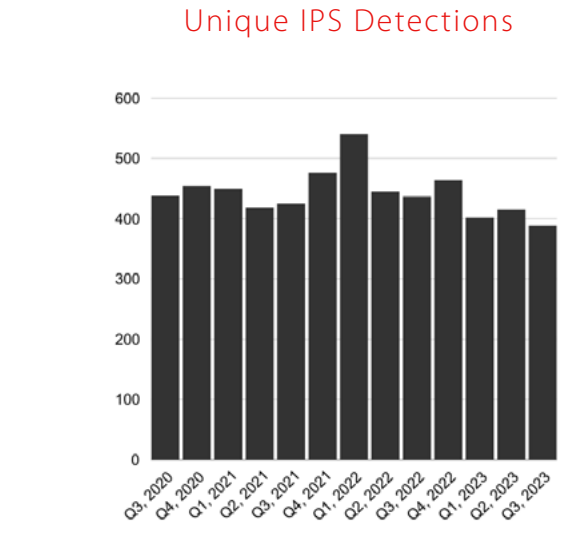

Figure 12. Unique IPS Detections

Additional Data:

- Total detections between Q3 2022 and Q3 2023 decreased by 343.42%. A less drastic number, we had a 1.24% average increase in total volume since Q3 2020.
- 3.42% increase in participating Fireboxes since last quarter.
- Q3 2022 to Q3 2023, there was a 12.34% decrease in unique signatures.
- There is a 5.3 average of new signatures in the Top 50 since Q3 2021.
- 8.20% of all the volume was for the ProxyLogon exploit (Signature 1138800).

Top 10 Network Attacks Review

The Top 10 signatures are the most voluminous among our 389 unique signatures this quarter. Many of the signatures represent popular vulnerabilities being exploited globally. They also show how old vulnerabilities continue to find popularity.

The top signature, 1138800, held its place for a second quarter in a row. It has been present in the Top 10 since Q3 2022. This is a Microsoft Exchange critical vulnerability known by the name ProxyLogon. While in 1st place both this and last quarter, its significance as the signature rose from 2.10% to 8.20% of total detections this quarter. The percentage per Top 10 signatures is notable because last quarter no signature surpassed 2.10%, with the lowest consisting of 0.70% of total detections. That balance is more so an outlier, as in previous quarters we commonly saw the Top 10 signatures ranging between 2-15%. In addition, the trend has been less top heavy. For a time, some top 3–5 signatures would range between 10-30% of total detections, per signature. Those extreme numbers continue to level out. We suspect this was due to some irregular Firebox traffic, even if it was still within the boundaries of statistical relevance. Since we updated our methods for processing this data last quarter, we expect to continue to see an increasingly balanced weighting among the signatures. That doesn't mean we won't see examples where signature 1138800 accounts for over 8% of total detections, but it seems the time when a signature representing 15-30% of total detections is over.

There are four new signatures in the Top 10 this quarter. Of the four, three have never been among the Top 50 signatures, looking several years back. The other five of six signatures were all present in the Top 10 last quarter. The sixth, signature 1058077, was present in the Top 10 in Q1 2023. In Q4 2022 it was the top signature. A SQL injection-type attack, the signature comprises various CVEs. Most notable affecting SCADA-based web-access portal software.

Signature 1132793, from the Top 10 signatures last quarter, moved up three places to the 3rd spot. It is for a vulnerability in ATutor software, an open-source learning management system (LMS). It is quite peculiar to see this not only remain in the Top 10 but rise further from the previous quarter. A no-longer-maintained education software is a very specific target. It makes sense to focus efforts against software likely riddled with vulnerabilities from the most recent years. What is surprising is seeing so many detections for software with a likely small community of users. Additionally, it's difficult to imagine any user or organization who is maintaining this software as being a lucrative target. This may be due to automated scanners poking around open ports, but even so, we would expect the numbers to result in a signature in the Top 10 that connected to a more widely used software. That is the case for many of the signatures, such as a few related to Microsoft-based technologies. We'll wait and see if this is a medium-term blip, or if the signature continues to maintain a large portion of total reported Firebox IPS detections.

Signature 1056161
While a completely new signature, it is for a 2012 vulnerability. CVE-2012-2329 is a buffer overflow vulnerability in the PHP's Common Gateway Interface (CGI) handling of web request headers. It will result in a buffer overflow and therefore allow attackers to initiate a denial-of-service attack. A successful attack against this vulnerability gives attackers code execution on the underlying web server, which some have actively exploited to place backdoors on web servers. The CVE connected to the signature affected PHP versions before 5.4.3. Now this is an outdated version, with PHP 8.3 set to be released at the end of November 2023. Therefore, we hope our customers are not using such outdated and vulnerable software. Regardless, the Intrusion Prevention Service should prevent known attacks like this.

Signature 1132401
SP2, 3.5, 3.5.1, 4.5.2, 4.6, and 4.6.1. Microsoft published the vulnerability along with a second CVE back in 2016, both discovered by Microsoft and neither under known exploitation until relatively recently. An attacker could potentially exploit the .NET Framework by sending an Extensible Stylesheet Language Transformations (XSLT) value to the server (via an XML client-side value) intending to cause a denial-of-service attack. The framework would try to compile the XSLT value recursively when transforming the content from XML.

Signature 1130660
This is a 2014 SQL injection vulnerability in Drupal version 7.x and any version before 7.32. The CVE, CVE-2014-3704, is a critical level vulnerability. This was discovered by researcher Stefan Horst, who posted a Proof-of-Concept a month after the Drupal team pushed out an advisory and update. This is considered a critical vulnerability because attackers could remotely exploit vulnerable Drupal deployments without any need for authentication. The seriousness of the vulnerability earned it the nickname Drupalgeddon. It is interchangeably referred to as Drupageddon without an "L", with Drupalgeddon being the name of the diagnostic tool published by Drupal. The Drupageddon/Drupalgeddon name may be familiar to many of you. Even though this was from nine years ago, there was another critical vulnerability of the same nature in 2018 that was then dubbed Drupalgeddon2.

The Drupal team published a module that administrators could use to check if they had any backdoors present, or any exploit traces. The module creators mentioned that this tool was of very limited value, since a Drupalgeddon exploit could be left with no trace. In addition, they state, "If you're still checking an un-patched or un-updated Drupal 7 site that is accessible to the public for hacks today, there's a strong probability that your site is already compromised."(**Source**). That statement confirms how serious of a vulnerability this is. While Drupal's latest version is now on 10.x, the Drupal version 7.x is still supported, with the planned end-of-life date set for January 5, 2025. Therefore, IPS does play an important role for any organizations continuing to maintain Drupal 7.x versions.

The creators of the tool published a lengthy and amusing **flow chart** on for administrators to assess whether they are vulnerable. It then follows all the necessary steps of the recovery phase for those unfortunate enough to have been hacked.

Signature 1056247

This signature is new to the Top 10 but has been floating among the Top 50 going back to at least Q4 2020. This signature is interesting in that it doesn't detect a specific vulnerability but instead detects an attempted exploit against a wide range of code execution vulnerabilities. Specifically, it detects a shellcode NOP Sled that attackers use while trying to exploit a buffer overflow vulnerability. A NOP Sled is a piece of shell code that uses the built-in No-Operations (NOP) assembly instruction to pad space in memory leading down to an instruction that gives an attacker full control of the process execution flow. This allows attackers to be less accurate when trying to overwrite a specific memory location because as long as the process execution lands on a NOP instruction, it will "slide" down until it hits the attacker-controlled instruction pointer. Memory stacks are often randomized for security reasons; therefore, a larger NOP sled will increase the likelihood of landing on it.

Two of the primary vulnerability exploits this signature detects are a 2005 code execution vulnerability affecting Squid proxy and another from 2017 affecting the Quagga network routing software suite. This signature went from the 28th spot last quarter to 5th now. Quite a jump like that is significant. There has been a consistent theme among several of the top signatures in recent quarters, which is that many of the top signatures are from major software products such as Microsoft Server or other management software. Neither Squid proxy nor Quagga fit into that category. Either way, attackers are directing efforts at these kinds of software and others with a NOP Sled exploit opportunity.

The Squid proxy vulnerability is from 2005. Squid is a caching proxy that serves several purposes, such as caching commonly accessed web pages to reduce bandwidth usage in addition to access time, filtering-traffic, and other support mechanisms for HTTP/S and FTP (with other protocols supported too). The Web Cache Communication Protocol (WCCP) is one of the other protocols supported by Squid proxy. It is a Cisco content-routing protocol that was later added to other vendors' products to increase efficiency with Cisco routers and switches. CVE-2005-0211 is specifically a flaw in the WCCP protocol in Squid. Attackers could successfully cause a denial-of-service attack by sending larger-than-supported WCCP messages than the memory buffer was intended to handle and in theory gain code execution. WCCP is disabled by default for Squid proxy installations, therefore minimizing the potential impact of this vulnerability. Patches were published to resolve this issue. For most readers (hopefully all) using Squid proxy, the need to check if you are properly patched should be a non-issue since this is a 2005 vulnerability. That said, it doesn't hurt to check if you have the latest patch. One last note – we mentioned vendors integrating WCCP into their product. While Squid proxy was one target, and their WCCP code perhaps unique to their product, other vendors may have been vulnerable to a similar vulnerability. The bulk of these detections may be due to automated exploit tools seeking exposed WCCP-integrated products with old and unpatched software. The other explanation for the large volume of detections could be from the second product affected by this vulnerability.

Quagga is an open-source network routing software suite for Linux and Unix-like systems that supports some of the most common routing protocols such as OSPF, RIP, and BGP. The project ran its course, and developers still interested in an open-source routing software suite switched their efforts to **Free Range Routing** (FRRouting, or FRR), a fork of Quagga. The Quagga project support had its last GitHub commit sometime in 2018, while FRR got its start in 2017. Over time, popular Linux distributions such as Ubuntu phased out Quagga from their package installations to instead promote FFR in its place. The CVE associated with Quagga is CVE-2017-549. Quagga daemons with telnet CLI enabled took in unlimited string lengths as long as a new line was not entered. Therefore, the system would eventually run out of memory, or the daemon would fail. Quagga version 1.1.1. addressed this issue. The vulnerability didn't solely affect Quagga. As the FFR project was forked from Quagga not too long after the vulnerability was published, FFR's code was similarly vulnerable to this exploit. Luckily, this was before FFR's inaugural release several months later. FFR addressed the issue. Therefore, FFR users do not need to worry about this vulnerability as the beta version was never publicly deployed. The real risk remains for anyone still using Quagga versions 0.93 to 1.1.0. Additionally, the Telent interface is restricted by default to local access, so the impact of this vulnerability should be significantly reduced.

Signature	Type	Name	Affected OS	Percentage
1138800	Web Attacks	WEB Microsoft Exchange Server Remote Code Execution Vulnerability -6 (CVE-2021-26855)	Windows	8.20%
1059958	Web Attacks	WEB Directory Traversal -27	Windows	6.23%
1132793	Web Attacks	WEB SQL injection select from attempt -5.a	Windows, Linux, FreeBSD, Solaris, Other Unix, Mac OS	5.63%
1056161	Web Attacks	VULN PHP HTTP_X Header buffer overflow	Windows	5.58%
1056247	Access Control	SHELLCODE NOP Sled	ALL	4.08%
1058077	Web Attacks	WEB SQL injection attempt -1.b	Windows, Linux, FreeBSD, Solaris, Other Unix, Mac OS	3.60%
1059877	Access Control	WEB Directory Traversal -8	Windows, Linux, FreeBSD, Solaris, Other Unix	3.37%
1132092	Buffer Overflow	FILE Invalid XML Version -2	Windows	3.36%
1132401	Access Control	WEB Microsoft .NET Framework Multiple Vulnerabilities (CVE-2007-0042)	Windows	3.33%
1130660	Web Attacks	DB Drupal Core database.inc expandArguments SQL Injection -3 (CVE-2014-3704)	Windows, Linux, FreeBSD, Solaris, Mac OS	3.32%

Figure 13. Top 10 Network Attacks by Volume

Signature 1138800 – ProxyLogon

The Microsoft Exchange vulnerability (ProxyLogon) is 1st in the Top 10 and 5th in the most-widespread. Additionally, it has been in the list of most-widespread since Q4 2022, except for last quarter. That the signature for ProxyLogon is back in the most-widespread list is expected of such a serious vulnerability. Disappearing from the most-widespread signatures last quarter isn't surprising as there was also a big decrease in total detections last quarter even though it reached the 1st place among signatures by volume. We have documented the history of signature 1138800 in several past ISR reports to demonstrate the rise of its prominence among other attacks. That is again presented in Figure 14.

Quarter	Rank by Volume	% of Total Volume
Q3 2023	#1	8.20%
Q2 2023	#1	2.10%
Q1 2023	#4	6.10%
Q4 2022	#4	5.54%
Q3 2022	#8	3.90%
Q2 2022	#14	1.80%
Q1 2022	#20	0.40%
Q4 2021	#26	0.30%
Q3 2021	#22	0.50%
Q2 2021	#20	0.60%

Figure 14. ProxyLogon history

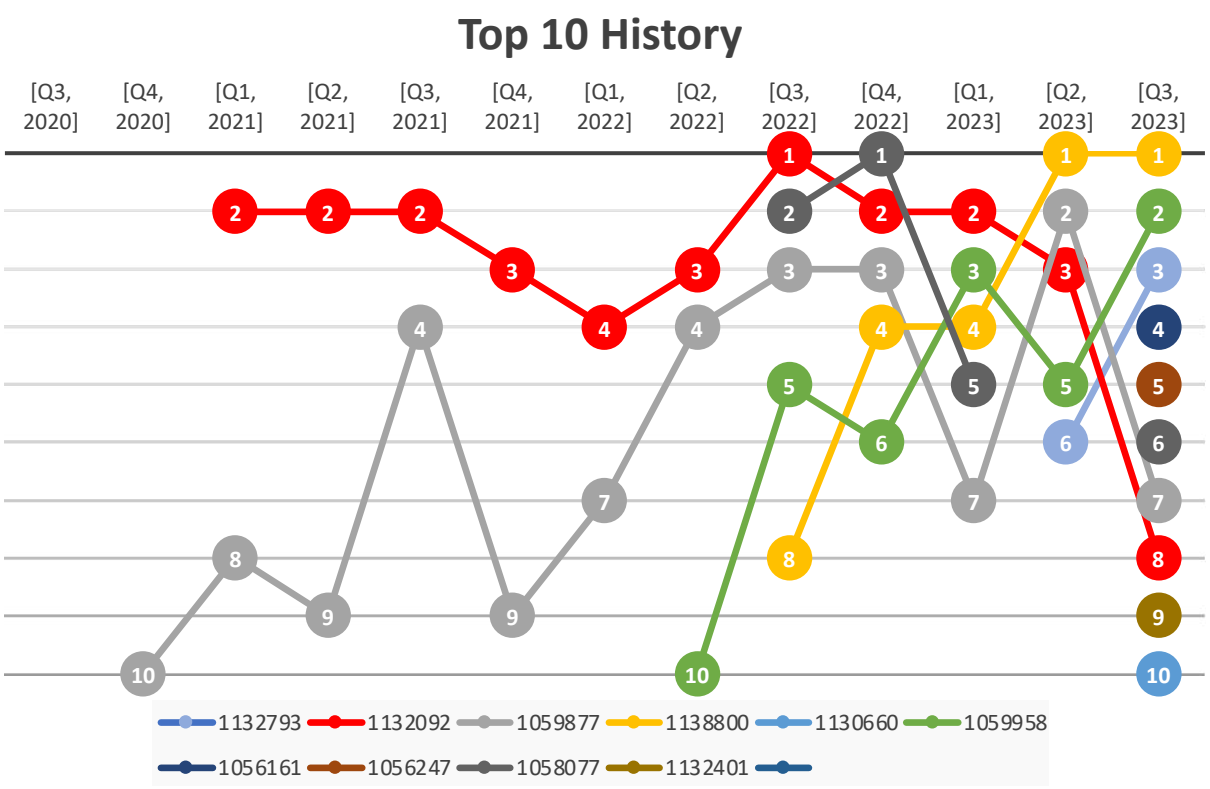

Figure 15. History of Prominent Signatures in the Top 10 Since Q3 2020

The Top 10 history chart displays the diverse patterns among our signatures. Several, such as signature 1138800 (yellow) and signature 1059958 (green), have had a progressive rise to the top within the context of the Top 10 signatures. That contrasts with signature 1132092 (orange) with a steady hold at the top but a quick drop since last quarter, as well as numerous other signatures that bounce around positions, but are ultimately on the decline. Then there are the four new signatures this quarter in 4th, 5th, 9th, and 10th place. On the far left of the X-axis is Q3 2020, with zero signatures present. There are only a handful of signatures that have remained in the Top 10 consistently, and often don't last more than two years. We might not see several of the long-term signatures in this chart come next quarter.

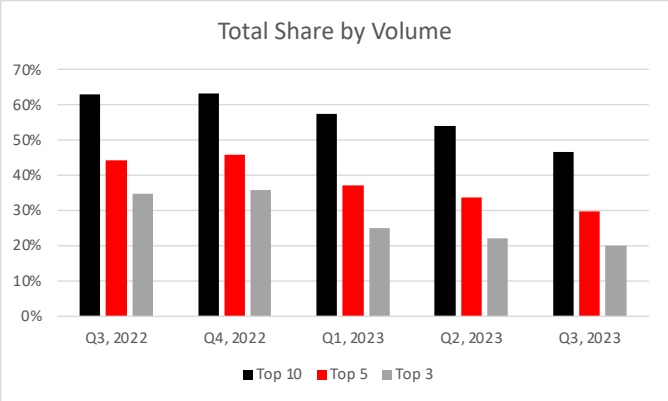

Figure 16. Total Share of Top Signatures by Volume Combined

We include Figure 16 to demonstrate how top-heavy signature composition is. The Top 10 signatures consist of 46.69% of total detections. The top 5 is 29.72% and the top 3 is 20.06%. A noticeable pattern is the continued downward shift in volume taken by the top trafficked signatures. Except for Q4 2022, increasing by a moderate amount. Even further back but not included in this chart, the Top 10 signatures had climbed from 72% in Q4 2020 to 86% in Q1 2022. Ever since then, the concentration among top signatures has continued to diminish. What does this all mean? One thought is that our data is simply improving. Many Fireboxes may have flooded the total count, even if they were within the statistical relevance of our processed data. That theory is semi-backed up as a few signatures that used to have a dominant place in our Top 10 are either no longer present in the Top 10, or if so, consist of a significantly smaller number of detections. Attackers commonly take advantage of new vulnerabilities, with "new" meaning a signature perhaps 3 years old or newer. That can be seen with the top signature, which is for ProxyLogon. It makes sense that attackers are directing their efforts to compromising high-value targets. That would steer a lot of total traffic to dominant signatures.

New Signatures in the Top 50

The IPS database is large, but it doesn't necessarily mean our customers will encounter each signature. This quarter, there were 389 unique signature detections (of those opted into data sharing). The comprehensive database is much larger than that. Among the unique detections we have tracked and discussed are the Top 10 signatures by volume. As of last quarter, we have begun identifying new signatures among our Top 50 signatures by volume. "New signature" in this context means the signature has not been present in the Top 50, but may have been present in the database for one or more quarters (or years). This might be a newly published signature, and it made it into the Top 50 this quarter as well. The three-signatures ranked above the 10th position have already been discussed, so we'll focus on the latter signatures.

Signature 1231674

In 13th place is a signature for the Microsoft Exchange Server server-side request forgery (SSRF) vulnerability (**CVE-2022-41040**), also known as ProxyNotShell. The name was given after researchers noticed the attack chain similarities to ProxyShell. Security researchers at GTSC **published** the zero-day on their blog on August 28, 2022 after privately reporting the vulnerabilities to Microsoft several weeks earlier. They decided to bring this forward to the public after noticing others falling victim to the attack. This affected Exchange 2013, 2016, and 2019. On September 29, 2022, Microsoft published an official CVE and **guidance** for the two vulnerabilities. At the time, they could only provide mitigation recommendations, of which one of the initial recommendations from Microsoft was immediately considered inadequate as security researchers identified easy workarounds. Microsoft soon updated that advice to ensure the URL Rewrite rule for Internet Information Services (IIS) server (one of many services running on Exchange) was less narrow in scope. In addition to adding this rule to IIS, Microsoft recommended disabling remote PowerShell for non-administrators. It wasn't until November 8th that Microsoft published a security update to address the pair of vulnerabilities.

Figure 17. Diagram of the Attack Using Two Vulnerabilities (Source)

How did this ProxyNotShell exploit find success? First, the attacker needed to be authenticated to the intended target's system. A successful password spraying attack or other initial access techniques were therefore necessary. Access to a non-administrative user's account would suffice for the attack, making the potency of the exploit greater. Next, the attacker used a specially crafted URL to perform an SSRF attack against the Exchange Autodiscover service to gain access to the privileged access of the Exchange PowerShell backend. They then used another publicly unknown vulnerability, CVE-2022-41082, to perform remote code execution using PowerShell. Additionally, CVE-2022-41040 allowed for remote triggering of CVE-2022-41082 (even though the attacker still needed initial access to the system).

Microsoft Exchange on-premises was vulnerable, while Exchange Online was not. That was not a relief for many as a large portion of organizations are using a hybrid approach to their Exchange utilization. Microsoft documents their observations prior to GTSC's public release of the vulnerabilities, where they noticed less than 10 organizations targeted. They attributed it (with medium certainty) to one state-sponsored organization, without naming them. There are suspicions it was a Chinese group based on IP addresses observed by GTSC, the use of the China Chopper web shell, and

several other Chinese attributes during the attacks. Once public, ProxyNotShell attacks took off, with many other groups integrating it into their arsenal of attack techniques. At any time, there are many publicly exposed Microsoft Exchanges servers, which is why a signature like this ends up 13th on our list. Attackers will maximize their opportunities.

Signature 1050435

This is quite an old vulnerability from 2006. The name "SHELLCODE Microsoft Windows CMD.EXE Reverse Shell" directly explains the attack. The detection of a Windows CMD.EXE banner via TCP is a possible sign that a malicious actor has achieved remote access (unless this was expected) through a spawned DOS command shell prompt over the TCP connection.

Signature 1133202

Last quarter we discussed a vulnerability for ATutor software that was tied to a new signature (1132793) in the Top 10. That signature is again in the Top 10, moving from 6th to 3rd place this quarter. We were a bit surprised to see an obscure and no longer maintained open-source learning management system (LMS) reach the Top 10. It's amusing that it is now the 3rd top signature by volume. Even more amusing is that there is another signature, 1133202, related to ATutor software. But after a quick review of both signatures, we see that this is for the same SQL injection vulnerability associated with CVE-2016-2555. Those with the affected 2.2.1 version can update to a later version. As we discussed in the last ISR report, even though users can address this vulnerability with an update, a lot of risk remains for anyone continuing to use this software. This is because it had not been updated for over four years.

Signature	Type	Name	Affected OS	Rank
1056161	Web Attacks	VULN PHP HTTP_X Header buffer overflow	Windows	4
1132401	Access Control	WEB Microsoft .NET Framework Multiple Vulnerabilities (CVE-2007-0042)	Windows	9
1130660	Web Attacks	DB Drupal Core database.inc expandArguments SQL Injection -3 (CVE-2014-3704)	Windows, Linux, FreeBSD, Solaris, Mac OS	10
1050435	Web Attacks	WEB Microsoft Exchange EwsAutodiscoverProxyRequestHandler SSRF(CVE-2022-41040)	Windows	13
1050435	Access Control	SHELLCODE Microsoft Windows CMD.EXE Reverse Shell -1.1	Windows	18
1133202	Web Attacks	WEB SQL injection select from attempt -5.x	Windows, Linux, FreeBSD, Solaris, Other Unix, Mac OS	48

Figure 18. New Signatures This Quarter

Most-Widespread Network Attacks

The most-widespread network attacks are determined by the greatest numbers of unique Fireboxes encountering a signature. The Top 10 signatures that we already discussed skew toward the sheer number of detections while still removing outlier Fireboxes. Therefore, any WatchGuard customer reading this will likely have seen a detection among our most-widespread signatures than the Top 10. This is the case for signatures 1059877 and 1138800. The Microsoft Exchange vulnerability (ProxyLogon) history for the Top 10 can be seen in Figure 14 in the Top 10 section.

Signature 1138800 is not the only one that has had a continued presence in the most-widespread list. Signatures in 2nd to 4th place are respectively, 1130592, 1110932, and 1059877. Last quarter they were in the same order but one place higher. Our one new signature this quarter is signature 1131523, a Microsoft Internet Explorer vulnerability.

Signature 1131523

The one new most-widespread signature this quarter is a Microsoft Internet Explorer memory corruption vulnerability published in 2015. A successful attack could cause a denial-of-service against the victim. It was one of numerous Internet Explorer vulnerabilities published by Microsoft on the same day. Internet Explorer 6 through 11 were vulnerable depending on the CVE, with CVE-2015-2425 associated with this signature, only affecting Internet Explorer version 11. What stood out from the announcement is that of the 19 published vulnerabilities, only CVE-2015-2425 and one other CVE were publicly disclosed, as well as it being the only one known to been exploited. The Microsoft bulletin for this and the other vulnerabilities was categorized as critical, with most using Internet Explorer 11 listed as Critical except for Windows Server 2008 R2 for x64-based Systems Service Pack 1 and Windows Server 2012 R2 considered moderate.

Signature	Name	Top 3 Countries by %			AMER %	EMEA %	APAC %
1131523	WEB-CLIENT Microsoft Internet Explorer Memory Corruption Vulnerability -2 (CVE-2015-2425)	UK 52.04%	Germany 45.31%	USA 43.66%	41.12%	41.78%	29.73%
1130592	WEB Apache Struts Wildcard Matching OGNL Code Execution -5	Brazil 39.13%	France 34.93%	USA 26.15%	25.21%	16.19%	10.81%
1110932	FILE Microsoft Windows GDIplus PNG tEXt Chunk Processing Integer Overflow	Portugal 21.3%	UK 21.06%	Brazil 19.25%	10.09%	16.96%	16.99%
1059877	WEB Directory Traversal -8	Germany 22.86%	Portugal 19.44%	Australia 16.83%	9.78%	16.07%	16.99%
1138800	WEB Microsoft Exchange Server Remote Code Execution Vulnerability -6 (CVE-2021-26855)	Germany 21.29%	Portugal 20.37%	Australia 12.87%	8.15%	14.14%	10.81%

Figure 19. Top Countries by %

The countries with the red-filled boxes in Figure 20 were listed as the most affected countries per the list of most-widespread attacks. In several cases, more than once. This list of ten countries, excluding Switzerland (has not been present since Q1 2022), have been the same countries one quarter or another among our most-widespread signatures. We chalk this up to two main reasons. One is that they are countries in a club of widely spoken languages. Attackers direct their campaigns where they can spread a wide net. Another is that they are all wealthy countries, or at least within the sphere of a middle-income country. Those reasons combined make them a reasonable target. Until a country doesn't meet both criteria, the logic will remain unchanged.

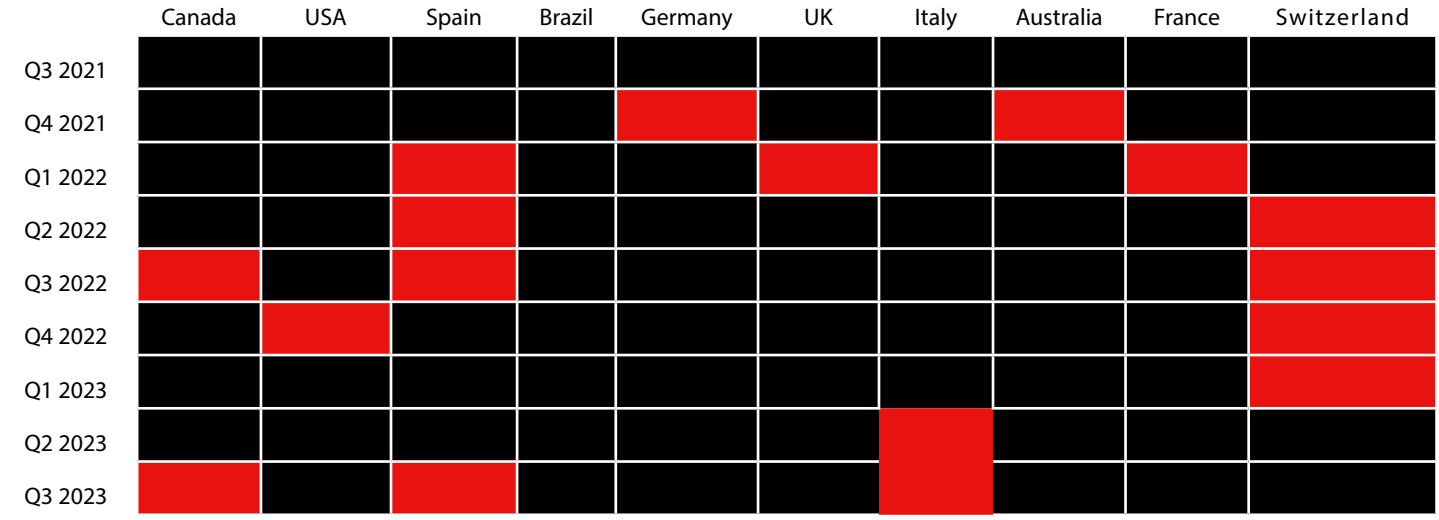

Figure 20. Countries Listed Among One or More Widespread Attack Signatures Who Were Most Affected

Network Attacks By Region

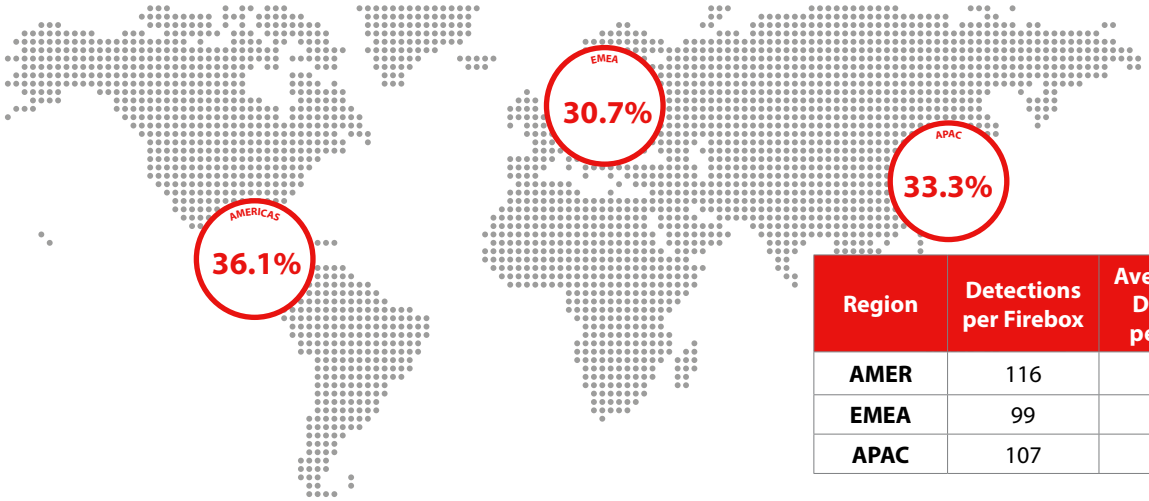

Region	Detections per Firebox	Average % IPS Detections per Firebox
AMER	116	36.05%
EMEA	99	30.66%
APAC	107	33.29%

WatchGuard customers are weightier in some regions than others. EMEA raw numbers are nearly double AMER, and AMER is nearly sixfold of APAC. Therefore, we normalize the data to show the true average detections per Firebox by region. We again see from last quarter that the total detection count is quite lower than what we've seen in the past several years. This is due to our change last quarter with how we changed the determination of relevant Firebox data. With that said, the data from this and last quarter are creating a new normal. Total detections per Firebox overall and regionally are climbing, but within a reasonable range from the numbers of the last quarter. That can be seen on the first chart in Figure 21. AMER went from 91 average detections to 116, and APAC from 67 to 107 this quarter. EMEA was relatively unchanged.

The second chart in Figure X shows the balance between the regions. The relatively equal balance between each region is something we haven't seen. Last quarter, AMER and EMEA were nearly even, but APAC had about a 10-point difference between them. This quarter, all regions are within a 6-point range. AMER stayed nearly identical to the last quarter, with a 0.02% shift.

There are a handful of factors that lead to regional shifts between quarters. One is regional cultures and their holiday schedule. That applies to both organizations and attackers. There are widely known patterns and inferences by the security community that hackers take holiday for several weeks or months during the summer, and holiday near the end of the year. But the holidays are also an opportunistic time for hackers to use their tool set. IT departments sometimes wind down major changes during the holiday and instead focus on maintaining a running environment on a lighter staffed schedule. Although the IT department may be working just as many long hours as before, it's an opportune time to do a regular audit of employees who are taking time off, or just catching up on work instead of putting out fires. With all these what-ifs for an organization's IT staff levels during holiday considered, none add "compromise of systems from ransomware or other attacks" to their weekly sprint plans. A well-prepared organization should have plans in place should this occur. A retainer for an incident response team is nice to have, but very spendy.

Another factor for shifting regional numbers may involve which customers enroll or unenroll in our telemetry sharing program. Often the noisiest Fireboxes comprise an outsized bulk of the total detection, even with statistical anomalies excluded. The numbers have decreased each quarter, but commonly the top 1% of Fireboxes represent over 40% of detections and the top 10% represent nearly 80% of detections. A third factor to consider is a change in patterns when high-volume Fireboxes in one region may then shift the balance presented in Figure 21. The numbers per region's are calculated in a way to prevent one region average from causing the imbalance to another, which is the reason for displaying detections per Firebox in Figure 22.

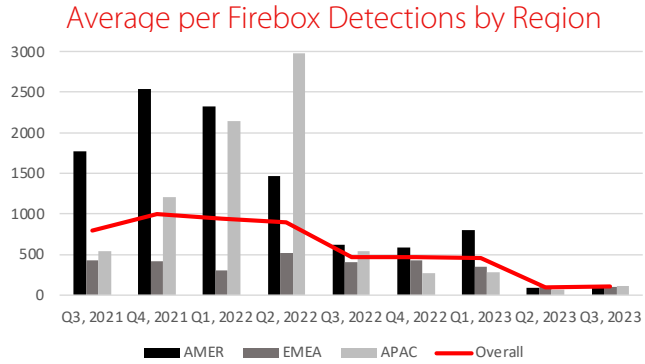

Figure 21. Average Detections per Firebox by Region since Q2 2021

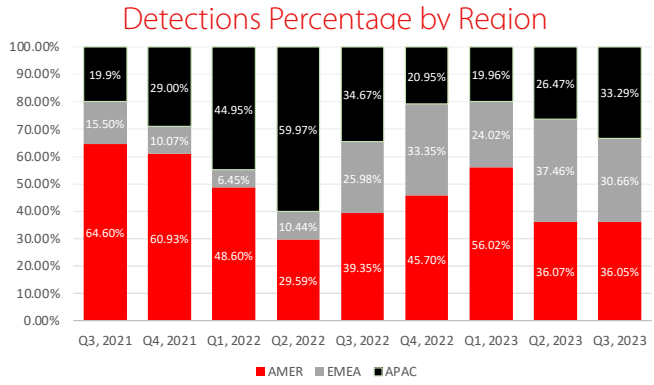

Figure 22. Average Detections per Firebox by Region

Conclusion

This has probably been hammered in, in some form or another, from previous conclusions in the IPS section, but if you are using a Microsoft product, you're going to forever need to be vigilant. The problem is that this goes for almost everyone. Concurrently, since many of us share this sentiment, we all know the drill. Patch your Microsoft Servers or have a bad time! Obviously, Microsoft isn't the only vendor with issues. That said, ProxyLogon is a pretty rough exploit. But other widely used products also create headaches for IT administrators. We saw a Drupal SQL injection vulnerability in the Top 10 this quarter. There are plenty of PHP and SQL-based web-facing systems that are in continuous need of patching. The best we can do, is do our best. That means implementing various tools to assist you such as vulnerability scanners, using good internal ticketing software, having system backups, and using defensive security tools to mitigate any attacks pre-patching.

DNS ANALYSIS

The modern Internet is built on using friendly human-readable names to access web resources instead of having to remember individual IP addresses of the service. Our software, whether it be a web browser or malware running on our machines, uses the domain name system (DNS) and DNS services to translate these human-readable names into network addresses. This makes DNS resolution an excellent place to identify and block threats before a network connection even establishes. In this section of the report, we analyze some of the top malicious domains that we protected WatchGuard DNSWatch customers from visiting in Q3 2023.

Top Malware Domains

Domain detections in this category include the websites attackers use to distribute malware and facilitate command and control communications. These domains typically do not have a legitimate purpose and are instead deployed specifically to enable malware infections.

Malware
x-vpn[.]ug
ocmtancmi2c4t[.]xyz *
thaus[.]top *
candatamsnd[.]info *
candatamsnc[.]info *
candatamsna[.]info *
candatamsnb[.]info *
t[.]hwqloan[.]com
xrass[.]com
carsfootyelo[.]com *

Figure 23. Top Malware Domains

WARNING

It should go without saying that you should not visit any of the malicious links we share in this report; at least not without knowing exactly what you are doing. Anytime you see us share a domain or URL where we have purposely added brackets around a dot (e.g. www[.]site[.]com), we are both making the hyperlink unclickable and warning you not to visit the malicious site in question. Please avoid these sites unless you are a fellow researcher who knows how to protect yourself.

There were seven new additions to the top malware domains by detection volume this quarter. We only just added the first new domain, ocmtancmi2c4t[.]xyz, in September 2023. Within only a month, it became the #2 malware domain by volume for the quarter. This domain is associated with HijackLoader, a popular first-stage malware that commonly loads information stealers like Luma.

We added the second new domain, thaus[.]top to the DNSWatch threat feed way back in December 2020 after finding it associated with command and control (C2) for the Phorpiex botnet. This botnet commonly intercepts cyptocurrency transactions and redirects funds to wallets under the attacker's control.

Figure 24. ViperSoftX PowerShell Loader

The next four domains, candatamsna[.]info through candatamsnd[.]info, are all related. We added these domains in August of this year, halfway through the quarter, after researchers found them associated with malware DNS Tunnelling efforts from a cryptocurrency mining malware variant.

The final new domain, carsfootyelo[.]com, also joined the threat feed late in the quarter after we found it delivering the IcedID malware variant. IcedID has been plaguing victims since 2017 as a second-stage remote access trojan commonly associated with Emotet.

Top Compromised Domains

Compromised domains may still house a legitimate website that an attacker has compromised to host malicious content. Attackers love to piggyback on the established good reputation of legitimate websites to host malware and phishing campaigns because they can often operate significantly longer without detection than if the attacker created their own purpose-built infrastructure.

Compromised
ssp[.]adriver[.]ru
d[.]zaix[.]ru
www[.]sharebutton[.]co
granerx[.]com
stopify[.]co
1[.]top4top[.]net
facebook[.]apps[.]fifyfive[.]co
www[.]cashconverters[.]sg
wieczniezywechoinki[.]pl *
dodgersdigest.com

Figure 25. Top Compromised Domains

There was only one new addition to the top compromised domains list this quarter, wieczniezywechoinki[.]pl. This domain houses a Polish Christmas decoration website that was compromised at one point to secretly host a malicious Office document masquerading as an invoice. If a victim opened the Office document, they were greeted with a request to select the "Enable editing" button, which then gave the document sufficient permissions to execute a macro and download the Emotet botnet.

Figure 26. dodgersdigest[.]com

Top Phishing Domains

As the category name suggests, detections categorized as phishing domains are websites we have found hosting phishing-related activity. Typically, these sites will mimic an authentication form for a legitimate web app like Microsoft 365 or Google Drive to trick victims into entering their credentials.

Phishing
unitednations-my[.]sharepoint[.]com
ulmoyc[.]com
edusoantwerpen-my[.]sharepoint[.]com
data[.]over-blog-kiwi[.]com
bestsports-stream[.]com
e[.]targito[.]com
t[.]go[.]rac[.]co[.]uk
www[.]898[.]tv*
nucor-my[.]sharepoint[.]com
googlestates[.]com*

Figure 27. Top Phishing Domains

There were two new additions to the top phishing domains list this quarter. The first new domain, www[.]898[.]tv, was associated with a Microsoft tech support scam. In this campaign, threat actors tried to trick victims into visiting a different domain microsoftassists[.]com which redirects to www[.]898[.]tv. This domain hosted a portable version of TeamViewer complete with a configuration file to automatically link it to the attacker's TeamViewer account, giving them immediate remote access to the victim's computer.

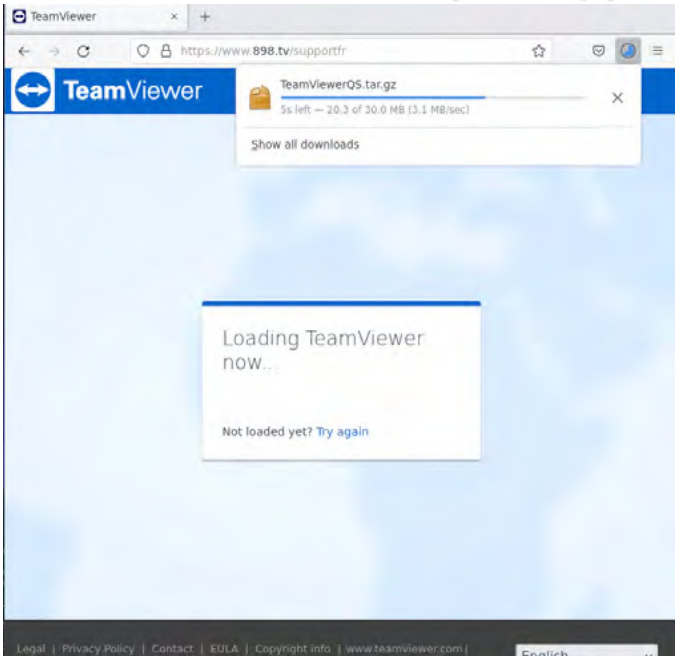

Figure 28. Fake Microsoft 365 Web App

The only other new domain to the top was googlestates[.]com. This domain came to us from a third-party feed after they found it associated with the JavaScript malware SocGholish. SocGholish is a "Fake Update" JavaScript malware that uses fake browser and application update notifications to trick victims into executing malware.

Conclusion

Over the last year, threat actors have pivoted to using legitimate remote access tools like TeamViewer as part of their attack campaigns. These attacks typically involve social engineering to trick victims into downloading and executing the software that then opens a backdoor for cybercriminals to swoop in. Traditional endpoint protection software is powerless to stop this style of campaign because the tools themselves are legitimate. This is why social engineering education and additional controls that can spot and stop phishing campaigns like DNS firewalling are so important.

FIREBOX FEED: DEFENSE LEARNINGS

Understanding what you are up against is an important start to building a layered defense. While the threat landscape is constantly evolving, there are a few trends that have solidified enough to be a consistent threat. Here are a few specific tips you can follow to address these threats.

01 Audit and Control Remote Access Software

This year has seen legitimate remote access software become a popular avenue for threat actors to gain access to an endpoint. By using legitimate software instead of a remote access trojan (RAT), threat actors can evade traditional anti-malware protections on the endpoint that by nature do not block goodware. One of the top threats in the DNS section of the Firebox Feed included one of these styles of attack. Administrators should recognize the shot across the bow and work to both audit existing remote access software used in their organization and block unauthorized use wherever possible.

02 Beware Malicious Email Attachments

Multiple top threats in the malware section of this report arrived through email messages either directly attached or delivered through a malicious link. While it can be difficult to block some attachment types, defenders absolutely should proactively block less-common attachment file types like .LNK and .HTA files. Organizations can take it a step further and consider blocking or at least quarantining document attachments from external sources and instead use authenticated file-sharing services like OneDrive to mitigate an entire attack vector.

03 Evasive Malware Is the Norm

While signature-based malware detection technologies still have their place in quickly identifying known threats, relying on these tools alone is a recipe for disaster. Motivated malware authors can iterate and improve their payloads faster than signature-based protections can keep up, meaning defenders must deploy advanced malware detection tools that use AI/ML and behavior analysis to stay safe.

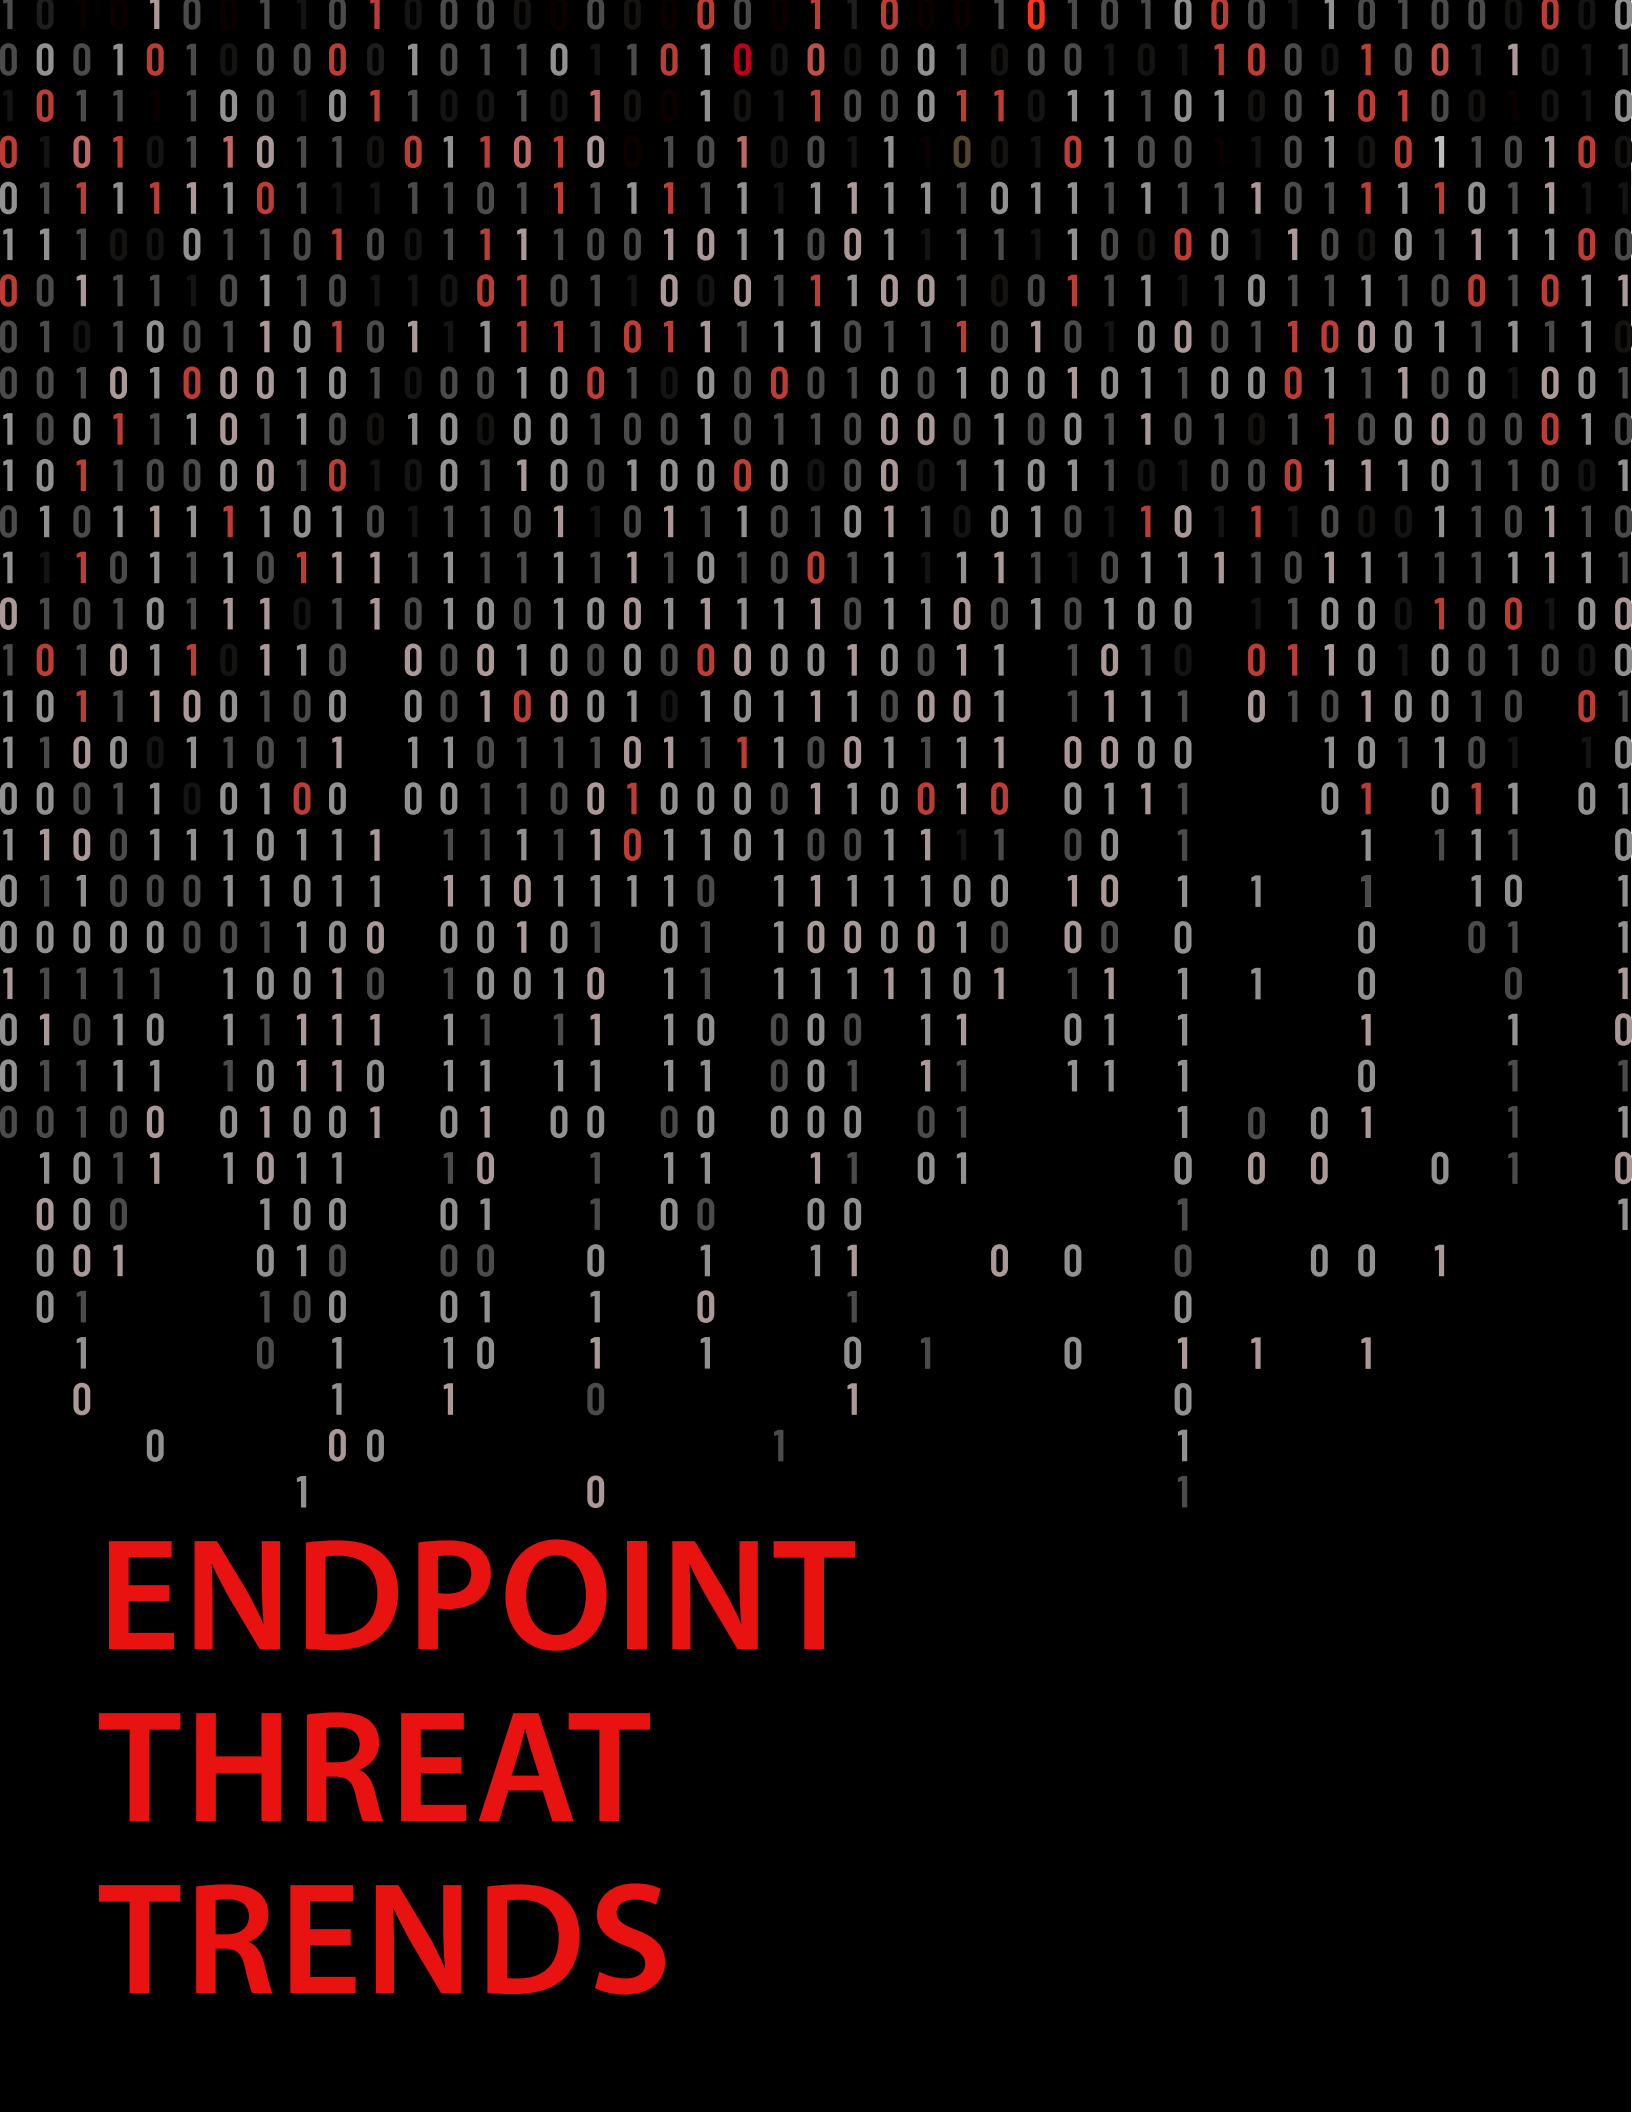
ENDPOINT
THREAT
TRENDS

If you've been following along for the past two iterations of the Internet Security Report, you'll know that we've revamped the entire report in the first quarter of this year and extensively included the Endpoint section in that endeavor. Previously, we included the total Firebox malware detections and broke those down into attack vectors, cryptominers, comparative analytics on browser-based detections, and the overall ransomware threat landscape. We've profoundly expanded on this data since incorporating the new changes in Q1. We still collect and highlight the overall malware frequency, attack vectors, browser-based detections, and the ransomware landscape. The only omission is cryptominers because those detections are associated with information-stealing malware.

The data we ingest and analyze is more extensive than ever before. Each quarter, we now collect:

- The total unique attacks blocked per 100k active machines
- The total observed malware hashes blocked per 100k active machines
- The number of alerts by the number of machines affected
- A ratio of the number of alerts over the number of machines for each country (alert coefficient), showing the top 30 affected countries each quarter
- The top 10 most prevalent malware
- The top 10 most prevalent Potentially Unwanted Programs (PUPs)
- The number of alerts by which WatchGuard technology invoked the alert
- Attack vectors
- Browser-based detections
- Alerts by exploit type
- (Threat hunting) MITRE ATT&CK tactics and techniques
- Firebox ransomware detections
- Ransomware group double extortions
- Notable ransomware breaches

In Q2, we discussed the lack of ascertaining patterns in the data because two data points (Q1 and Q2) weren't enough to make any determinations. Well, it's Q3, and we've upheld our promise of ensuring we share these patterns. Almost every subsection hereinafter highlights this quarter's data, comparative Q2 data, and 2023 quarter-over-quarter (QoQ) data. In addition to the recurring data collection practices, we've also improved many of the graphs and tables from the previous quarter and even added a few never-before-seen ones in the Malware Frequency, Attack Vectors, Threat Hunting, and Ransomware Extortion Groups subsections.

We're thankful for the users who opt in to our anonymous data collection program via the Firebox, and to show our appreciation, we have continuously made enhancements to our analysis to give back. By leveraging WatchGuard's Endpoint Protection, Detection and Response (EPDR) solution, we block the latest threats on every applicable endpoint, extract the telemetry from that attack, and put it through different lenses, to show how EPDR protects your data and your organization. The subsections below show you these lenses and our findings, beginning with Malware Frequency, as usual.

MALWARE FREQUENCY

As a quick refresher, we denote malware frequency as "per 100k active machines." In other words, if your organization has 100,000 machines with an active EPDR license, the number below is the number of unique attacks you would have blocked. In Q3, we saw a steep decline in the number of unique attacks blocked per 100k active machines at 171, a reduction of 82.57% from Q2. However, before we continue, we want to provide some context for this data we hadn't stated prior. This number is for unique attacks blocked per quarter, which doesn't describe the total number of attacks blocked, including previously known malware. A unique attack is a malicious file we haven't seen before; it's a unique hash.

Unique Attacks Blocked per 100k Active Machines	171

There's a clear pattern here – we are observing fewer unique malware detections QoQ, especially in Q3. Is it that WatchGuard's EPDR is detecting fewer of these attacks? Emphatically, no. Our data suggests we observed similar, if not more, overall indicators of compromise (IoCs). A clear example of this lies within our Threat Hunting data later in this section. We observed almost three times the number of threat-hunting alerts from the quarter prior.

Furthermore, our EPDR data set tells us that we averaged a little over 550,000 IoCs per day across slightly under 80,000 machines—almost seven IoCs per machine per day (~6.875). For reference, in Q2, we observed an average of around 487,000 IoCs per day across 73,000 machines for an average of about 6.67 IoCs per machine daily. Thus, we saw more IoCs across more machines in Q3. However, we observed most IoCs from already known, documented malware.

We know of two reasons for some, if not most, of the frequency reduction described before; one is the takedown of QakBot (Qbot) in the middle of Q3 by law enforcement from several countries. The threat actors behind Qbot campaigns were responsible for many of the daily IOCs and malware alerts. The email campaigns were relentless, ever-changing, and widespread – the perfect ingredients for a swarm of new and existing IoCs.

The other reason is that three out of the top four Top 10 Malware for this quarter were on the list last quarter (described later). This finding means every alert invocation from those three malware didn't count toward the Unique Attacks Blocked per 100k Active Machines counter above. Compounding the fact that the malware

frequency was previously trending down from Q1 to Q2, this results in a significant reduction in overall malware frequency. Fortunately, we are introducing another data point this quarter that

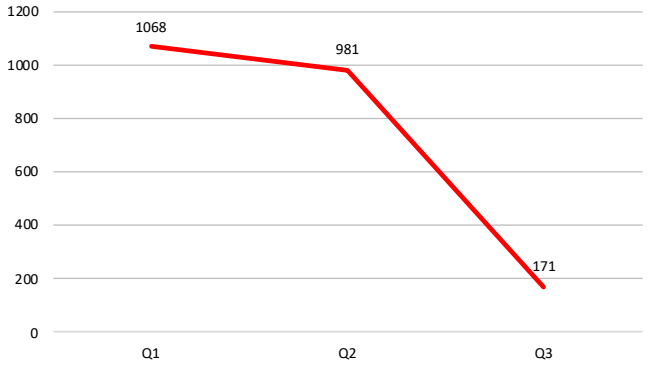

Figure 29. 2023 QoQ Unique Attacks Blocked per 100k Active Machines

provides additional malware frequency telemetry, showing similar malware frequency from Q2 – the total number of unique malware hashes we observed. We still denote this number in the "per 100k active machines" unit of measurement. For Q3, WatchGuard EPDR blocked 1.09 malware hashes per 100k active machines. Another way of thinking about this data point is that the 1.09 number means that for an organization with 100,000 endpoints, about one malware attack would have penetrated your network if it were not for EPDR. Furthermore, since we are just now introducing this data point in Q3, we also decided to include Q1 and Q2. You can see them in the figure below.

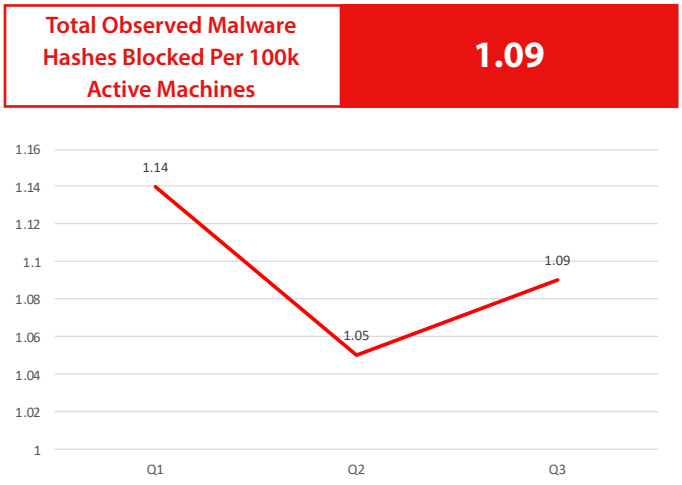

Figure 30. 2023 QoQ Total Observed Malware Hashes Blocked per 100k Actve Machines

Alerts by Number of Machines Affected

This subsection describes the data through the lens of determining how many machines an alert appeared on. By knowing which alerts appeared on only one machine versus those on hundreds or thousands of machines, we get a glimpse of whether threat actors are performing more precise attacks or spamming email phishing campaigns, like Qbot described previously. Using Qbot as another example, this quarter, we saw a curtailment of alerts appearing on over 100 machines (-24.05%), tangible proof that the Qbot infrastructure takedown impacted our data.

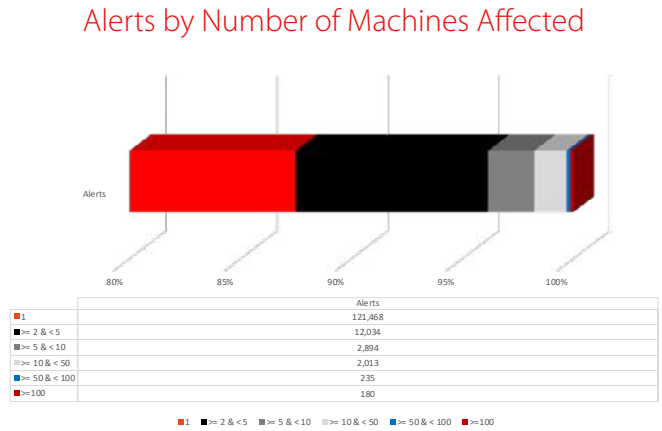

Figure 31. Alerts by Number of Machines Affected

The bullet points below define and describe the parameters for which we log this data:

- **1** – Exactly one machine alerted on this file/process.
- **>=2 & < 5** – Between two and five machines alerted on this file/process.
- **>=5 & < 10** – Between five and ten machines alerted on this file/process.
- **>=10 & < 50** – Between ten and fifty machines alerted on this file/process.
- **>=50 & < 100** – Between fifty and 100 machines alerted on this file/process.
- **>=100** – More than 100 machines alerted on this file/process.

In addition to the reduction in alerts on more than 100 machines, alerts on those between two and five (-18.58%), five and ten (-5.95%), and ten and 50 (-4.55%) also decreased from Q2. The only parameters that increased from Q2 to Q3 were those alerts on only one machine (1.45%) and those between 50 and 100 machines (6.33%). However, these were only modest increases. When accounting for volume, the number of alerts practically nullifies their differences for this quarter. Ultimately, this quarter's data shows a modest rise in single endpoint attacks and a reduction in widespread campaigns.

Number of Machines	Q2 Alerts	Q3 Alerts	Difference from Q2	Percentage Difference from Q2
1	119,735	121,468	1,733	1.45%
>= 2 & < 5	14,781	12,034	-2,747	-18.58%
>= 5 & < 10	3,077	2,894	-183	-5.95%
>= 10 & < 50	2,109	2,013	-96	-4.55%
>= 50 & < 100	221	235	14	6.33%
>=100	237	180	-57	-24.05%

Figure 32. Alerts by Number of Machines Affected

Alerts by Top 30 Countries Affected

This subsection is interesting as it looks through the data with a geographical lens. In other words, we show the top countries affected by malware this quarter. If we were to take the total number of alerts and rank the leading countries, that would be misleading because those countries with the most users would consistently rank at or near the top. Therefore, we defined the Alert Coefficient variable, a ratio of the total number of malware alerts over the number of active endpoints for each country. This variable ensures we more accurately represent countries with few WatchGuard EPDR-protected endpoints and those with many more endpoints.

Let's take Malawi as an example to explain the Alert Coefficient. Our data shows us that Malawi had an Alert Coefficient of 1.00. Let's say that 100 users in Malwai had WatchGuard EPDR and opted into our anonymous data reporting program. For Malawi to have an Alert Coefficient of 1.00, there would have to be 100 total alerts in Q3 (100 alerts / 100 endpoints). However, even though Malawi did have an Alert Coefficient of 1.00, the 100 endpoints and alert numbers were fictional. Ironically, Malawi had the second-highest Alert Coefficient but had the 30th most total alerts.

Five new countries appeared in the top 30 that didn't appear last quarter – New Caledonia, Nigeria, Zimbabwe, Singapore, and Trinidad and Tobago. New Caledonia is surprising as it was new and appeared at the number one rank with a record-high 3.70 Alert Coefficient. This Alert Coefficient means that for those users in that country, there were, on average, 3.70 detections for each machine. The other four countries appeared near the bottom of the list with Alert Coefficients between 0.11 and 0.07.

The rest of the reappearing countries were a bit of a mixed bag. Many of the countries moved one or two spots up or down. However, there were a few exceptions. For example, Armenia saw the most significant leap in the rankings, going from last place in Q2 to 16th place this quarter. Bosnia and Herzegovina and Paraguay also saw large rank increases, moving up seven spots each. On the contrary, the Cayman Islands saw the most significant decrease in the rankings, moving down 12 places from the quarter prior. An honorable mention for those moving down in the rankings was Turkey, which repelled down seven spots from the last quarter to the 26th spot.

Country	Alert Coefficient	Order Difference from Q2
New Caledonia (France)	3.70	*NEW*
Malawi	1.00	*+1*
Laos	0.85	*+1*
Jordan	0.81	*-2*
Cuba	0.75	*-4*
Morocco	0.58	*+1*
Pakistan	0.54	*-2*
Bosnia and Herzegovina	0.44	*+7*
Mozambique	0.36	*+1*
Angola	0.33	*-1*
Vietnam	0.31	*+3*
Sao Tome and Principe	0.25	*-1*
Kenya	0.19	-
Bolivia	0.19	*+2*
Bangladesh	0.17	*+2*
Armenia	0.16	*+14*
Macedonia	0.15	*+4*
Paraguay	0.13	*+7*
Botswana	0.13	*+5*
Cayman Islands	0.12	*-12*
India	0.12	*-1*
Guatemala	0.11	-
Nigeria	0.11	*NEW*
United Arab Emirates	0.11	*-1*
Zimbabwe	0.10	*NEW*
Turkey	0.09	*-7*
Venezuela	0.09	*+2*
Indonesia	0.09	-
Singapore	0.08	*NEW*
Trinidad and Tobago	0.07	*NEW*

Figure 33. Alerts by Top 30 Countries Affected (with QoQ Differences)

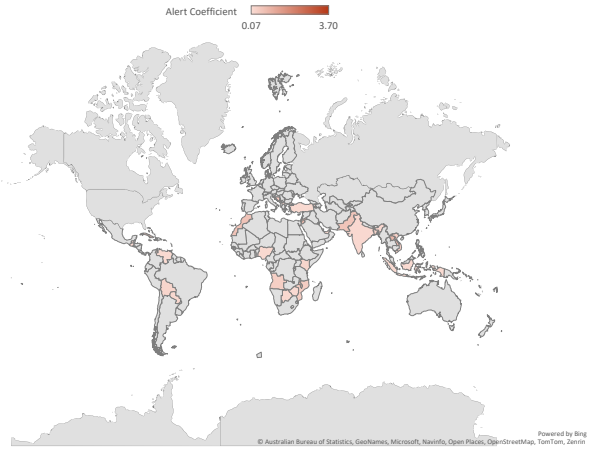

Figure 34. Alerts by Top 30 Countries Affected (Map)

TOP MALWARE AND PUPS

The most prevalent malware and PUPs are a favorite amongst readers, providing insight into the most ubiquitous hashes seen each quarter. This data is helpful for both WatchGuard and our users because it allows us to take extra precautions to protect against these specific threats, knowing that these are the most widespread campaigns in the wild. Furthermore, in a way, it also shows which malicious email campaigns are successfully tricking users into clicking links or downloading attachments. For example, the Top 10 Most Prevalent Malware list often contains Glupteba and GuLoader – malware known to spread via phishing campaigns – and this quarter is no exception. Threat actors commonly use phishing emails as their entry point into a network. If those emails successfully trick the user, WatchGuard EPDR blocks their execution, and we log it.

Top 10 Most Prevalent Malware

Speaking of Glupteba, it was the most prevalent malware this quarter, and GuLoader appeared on the list twice at ranks eight and nine. GuLoader was the only malware appearing more than once this quarter, and both Glupteba and GuLoader have appeared in the Top 10 Most Prevalent Malware list every quarter since we began logging this data. In Q2, the reappearing Glupteba sample ranked second before claiming the number one spot this quarter.

Two other hashes appeared this quarter and last: MyloBot and an unknown malware (injector) not attributed to a specific family. In Q2, this specific MyloBot variant ranked sixth, but in Q3, it moved up three spots to third rank. On the other hand, the unknown malware injector remained in the number four spot. We now denote recurring hashes in the top 10 list from the prior quarter with an asterisk (*).

To our surprise, Agent Tesla finally made the Top 10 Most Prevalent Malware list for the first time, and just barely, ranking tenth. Anecdotally, we observe a lot of Agent Tesla samples, but they all have different hashes. This anecdote means Agent Tesla is possibly more prevalent than the list gives it credit for. Just because a single hash is the most prevalent for any quarter doesn't mean those malware families are the most observed. It means that specific variants of these families are more effective or pervasive than others. In other words, the Top 10 Most Prevalent Malware list shows unique hash prevalence instead of malware family prevalence.

Two other malware families appeared on the list for the first time – Valyria and Medusa ransomware. Valyria, like many other malware campaigns, is disseminated via phishing emails. This quarter, the Valyria sample in the list leveraged a malicious macro-embedded Excel spreadsheet called "declaracion_de_transaccion.xls." The embedded macro attempts to download additional malware once the user downloads and executes this file. However, if the user had WatchGuard EPDR, this execution was blocked. In Q3, we stopped 1,158 instances of this Valyria malware campaign.

The Medusa ransomware variant appearing in the list is also a surprise, especially considering that the ransomware operators specially crafted this sample for the victim organization, which we will not name. Since 460 machines logged this block, it infers the threat actors attempted to encrypt 460 machines on this organization's network, and EPDR blocked all of them. We are exceptionally proud of this defensive action because it proves that EPDR successfully stopped known malware at scale and saved this organization from an imminent disaster.

Glupteba

Glupteba is a multi-faceted malware-as-a-service (MaaS) with capabilities such as (down)loading other malware, acting as a botnet, stealing information, stealthily mining cryptocurrency, and more that targets victims seemingly indiscriminately worldwide. In 2021, Google disrupted the botnet, but it made a resurgence in late 2022 into early 2023. Like GuLoader, threat actors commonly use evasive downloaders to deliver additional malware. Although, unlike GuLoader, Glupteba is arguably more sophisticated and has more capabilities. It's an evasive trojan that researchers have observed taking control commands from the Bitcoin blockchain, among many other techniques for evasion.

Valyria

Threat actors disseminate Valyria malware almost exclusively via email phishing and spam campaigns. They almost exclusively use Microsoft Office solutions such as Word and Excel. To make the emails seem more believable to unsuspecting victims, the threat actors will password-protect these attachments and provide that password in the email body. The admission of a password makes it seem more realistic and from an actual, non-malicious entity. However, when the user enters the correct password after opening the Microsoft-related document, malicious embedded macros perform tasks behind the scenes, usually downloading other malware. The Valyria example in our Top 10 list leveraged an Excel file that targeted Spanish-speaking users.

MD5	Signature	Affected Machines per 100k	Classification Attestation
6CC8D5F1CB1819791E4897F902FAF365*	Trj/RnkBend.A	1,262	Glupteba
AE484F4E3FA0415F62DDDE614D8E30BE	W97M/Downloader.DDE	1,158	Valyria
3E86685246C1FDCC9EEF8B95986BA4E4*	Trj/WLT.F	634	MyloBot Delivering Khalesi
2253836BB8B0B5479A1F77974B82B1F0*	Trj/RnkBend.A	538	Unknown Malware (Injector)
1A527D9250D86BE6759FCE3FAD7093FF	Trj/Agent.JTM	469	Unknown Malware (Dropper)
E0485EA9057387DDFE8C2272FDB01333	Trj/GdSda.A	460	Medusa Ransomware
05B8D66F4856D0161B01A1FC29037AA2	Trj/CI.A	310	Trojanized VLC Torrent Installer (Shellcode)
5A748796698A5C76D231512B2426A231	Trj/Agent.MK	280	GuLoader
8D5332901CB81F3BC447FD81324E06FE	Trj/Chgt.AD	226	GuLoader
79CDE4ABBFEF7669F6BCA336E2BB6D20	Trj/Chgt.AD	221	Agent Tesla

Figure 35. Top 10 Most Prevalent Malware

MyloBot
MyloBot has been active for around five years, and interestingly, the botnet operators are known to have attempted to extort victims via email. More ubiquitously, the malware's primary intent is to infect a machine without the victim's knowledge, allowing attackers to leverage any device within its botnet to perform actions on the attacker's behalf. Like other botnets and loaders, the malware downloads the final payload after multiple stages of evasively downloading malicious files in a daisy-chain fashion.

Unknown Malware (Injector)
An "unknown malware" is one we can't attribute to a specific malware family, but we can at least generically identify it as a malware tool. An injector is malware that injects itself or a payload into another process. An example is when malware creates a process in suspended mode, injects a payload into it, and continues its execution.

Unknown Malware (Dropper)
This malware is a sample we cannot directly attribute to a particular family but generically identify it as malicious software. In this case, a dropper is a malware that "drops" another malware, as the name suggests. An example of a dropper is an embedded payload that is de-obfuscated at run time and placed on the victim's machine.

Medusa Ransomware
The Medusa ransomware is from the group of the same name – Medusa. They operate a dark web data leak site called Medusa Blog, where they post all the victims unwilling to pay a ransom in a tactic called "double extortion." The payload, sometimes called the encryptor, is the file that ultimately encrypts all files on a system and drops a ransom note. That is the file that appeared in our Top 10 this quarter. The group creates semi-custom ransom notes identifying the organization they intend to encrypt. We will not name that organization in this report. However, we will say that the payload uses AES256 encryption coupled with an RSA public key to encrypt files. It then appends all encrypted files with a ".MEDUSA" file extension. Finally, we call this ransomware human-operated ransomware (HumOR) because it allows command line parameters, meaning a human must operate it. Although, they could also write a script that automates the final payload. Nonetheless, it's still human-operated.

Trojanized VLC Torrent Installer (Shellcode)
This sample on our top 10 list appeared to be a genuine VLC Torrent Installer. However, the file contains embedded shellcode that performs malicious actions when executing. Users who run this file will experience an installation wizard that appears to be legitimate, and it does successfully download VLC Torrent software, but, as has been stated, it also performs malicious actions in the background. To the layman, they will have no idea that any malicious action occurred. However, WatchGuard EPDR blocks this and alerts the user to the malicious intent.

GuLoader
This malware is sent in waves by attackers who send out spam phishing emails with malicious attachments containing the first stage of their campaigns – GuLoader. GuLoader is commonly used to download additional malware, such as infamous information stealers like RedLine Stealer, Racoon Stealer, Vidar, and FormBook. It is persistently on the top 10 list, or close to it, and is the most observed prevalent malware since we've started tracking this data.

Agent Tesla
Agent Tesla is another information stealer and remote access trojan (RAT). It's been one of the most prevalent for the past several quarters. Surprisingly, it made the top 10 list for the first time in Q3 because there are a lot of different versions. It's difficult for one single hash to affect so many machines as opposed to other spam malware campaigns such as GuLoader and Glupteba. Agent Tesla is a .NET program that appears to be an authentic file. These files come in various types, but threat actors fully coded them to appear as authentic as possible, appearing as calculators, educational programs, and more.

Top 10 Most Prevalent PUPs

A PUP is an acronym for a potentially unwanted program. You may also commonly see these as PUAs or potentially unwanted applications. These are the same thing. The common denominator is that these terms describe software that is unwanted or acts suspiciously based on the context of the file. For example, AutoKMS files are PUPs because they allow users to circumvent Microsoft license agreements illegally. These files aren't malicious – they won't damage your computer – but they are programs users might not want to be associated with. They are potentially unwanted programs/applications.

This quarter, there were four AutoKMS PUPs, three BundleOffers, two Hacktools, and one PortScanner. All AutoKMS-related files in the top 10 perform similar but slightly different actions. The number one rank, KMSPico, and the seventh-ranked file, AutoPico, are practically the same. They are both activators of illegal Windows products, particularly Windows operating systems. However, those who use KMSPico also leverage it for Microsoft Office products, hence why it is ranked first this quarter. Microsoft Toolkit, which was ranked fifth, performs similar actions. The other AutoKMS PUP that appeared in the list this quarter is a software installer, which we couldn't attribute to any specific software.

The three BundleOffer PUPs are all installation wizards of PDF-related software. Two are from PDF Power, and the other is from PDFCreator. We found that these files aren't malicious. Still, they include an installation wizard that tricks or coerces users into downloading additional software, usually adware or third-party software. Because users typically don't want this extra software, we label them PUPs.

The final three in the Top 10 Most Prevalent PUPs are different hacking tools. We labeled two of these as generic hacking tools; the other is a port scanner. Rank sixth is an extensive JavaScript file that exploits the Heartbleed bug (CVE-2014-0160) disclosed almost ten years ago. According to those who disclosed the bug, Heartbleed "allows anyone on the Internet to read the memory of the systems protected by the vulnerable versions of the OpenSSL software." Attackers exploit this bug by compromising the secret keys to identify the service providers that encrypt the network traffic. This compromise allows them to eavesdrop on traffic unabated.

The other generic hacking tool we couldn't attribute to any software or tool. It is likely a proprietary tool used by red teamers or hackers. Since both nefarious users and genuine penetration testers can use hacking tools, we label most of these hacking tools as PUPs. Given enough telemetry and context, we classify some hacking tools as malware, but not often. The final hacking tool, a port scanner, is ranked tenth and is the genuine software Advanced Port Scanner. As the name suggests, this tool scans ports of any given IP address.

MD5	Signature	Affected Machines per 100k	Classification Attestation
8D0C31D282CC9194791EA850041C6C45	HackingTool/ AutoKMS	12,750	KMSPico
30C7E8E918403B9247315249A8842CE5	HackingTool/ AutoKMS	9,527	Unknown Software Installer
FB396E6E8B08308F8D12F2776EDA4C85	PUP/BundleOffer	2,540	PDF Power 3.0.0.0 Setup Wizard
01C283988C93D390D4C81C38BF00ABEE	PUP/BundleOffer	1,771	PDFCreator 5.1.2 Setup Wizard
EE7714229183964C8AA1FC8FE0C8CEED	HackingTool/ AutoKMS	906	Microsoft Toolkit 2.6.4.0
1E2A99AE43D6365148D412B5DFEE0E1C	PUP/BundleOffer	765	PDF Power 4.0.1.0 Setup Wizard
C9E4916575FC95BEDBD12415AB55CC84	PUP/Hacktool	755	CVE-2014-0160 (Heartbleed) JavaScript Exploit Script
CFE1C391464C446099A5EB33276F6D57	HackingTool/ AutoKMS	748	AutoPico
CD8AF8E8A07D6C58A500A23B501560B6	PUP/Hacktool	746	Unknown Hacking Tool
6A58B52B184715583CDA792B56A0A1ED	Hacktool/ PortScanner	657	Advanced Port Scanner

Figure 36. Top 10 Most Prevalent PUPs

HackingTool/AutoKMS
AutoKMS is an umbrella term encompassing any cracked Microsoft software that allows users to use Microsoft products without a license, or it's a file that facilitates the bypass of Microsoft licensing.

PUP/BundleOffer
A classification reserved for installers that include third-party software or "offers." Usually, the third-party software is adware, which is particularly unwanted.

PUP/Hacktool
PUP/Hacktool is a generic classification for any tool or software used for hacking purposes. Both legitimate penetration testers and malicious threat actors use these tools. For this reason, we classify these as PUPs because we can't be sure whether these tools are malicious. However, if we capture telemetry or additional context that allows us to determine if a malicious threat actor uses a hack tool, there's a chance we classify it as malware. Most open-source tools are PUPs or goodware. It's the proprietary ones that we usually label as malware.

Hacktool/PortScanner
This signature is yet another generic classification for a hack tool, but with a bit more specificity. Hashes with this classification perform port scanning actions on networks. Like the PUP/Hacktool classification above, we can't be sure whether a penetration tester or malicious threat actor uses these tools. If given more information, we could make a more specific determination.

Defense in Depth

This subsection highlights the efficacy of WatchGuard's EPDR holistic solution. What we mean is that EPDR is a multi-faceted solution that acts as a technological phalanx on the endpoint that can detect, block, and prevent malware in various ways. The Defense in Depth subsection pulls back the curtains on EPDR by showing which technology blocked an attack that arrived on an endpoint. Knowing this information not only indicates which technology is responsible for the most blocks but also shows how a defense-in-depth solution is a posture recommended for any organization of any size.

EPDR comprises six technologies:

- **Endpoint Detection** – The typical, legacy endpoint antivirus solution, Endpoint Detection displays the number of hashes invoking an alert located in our known-malicious hash database. This is commonly called a signature-based detection antivirus solution.
- **Behavioral/Machine Learning** – Behavioral/Machine Learning is a step above signature-based detections because it analyzes the file's actions upon executing in a sandbox. We create rules based on these behaviors and determine whether they are malware.
- **Cloud** – Alerts that fall under the Cloud category are files sent to WatchGuard's Cloud servers for further analysis beyond signature-based detections and behavior/machine learning. The files that are malicious iterate the counter here.
- **Digital Signature** – Digital Signatures are methods of determining the authenticity and legitimacy of the sending user and ensuring it hasn't been tampered with (integrity). We make malware determinations based on these digital signatures. If an attacker altered it in transit, it is a digital signature from a known malicious user, or if we know the signature is compromised, we make a further decision.
- **Manual Attestation** – Manual Attestation is a fancy way of saying that a human analyst scrutinizes the file. If the file makes it past all of the other technologies and still looks suspicious, one of WatchGuard's attestation analysts performs the analysis and makes a classification. Once a file reaches this stage, a classification, whether goodware, PUP, or malware, is always determined.
- **Defined Rules** – The final technology, Defined Rules, are predefined behaviors that, if a file were to perform, we would determine are malware. Most people associate defined rules with threat hunting, but these rules can apply to endpoint detection, too.

In Q3, we saw a significant jump in Endpoint Detection alerts, increasing 29.36% from Q2. We also saw a massive increase in digital signature alerts, rising by a whopping 735.21%. Typically, Digital Signatures is the technology with the least number of alerts. However, it superseded Defined Rules and Manual Attestation this quarter, becoming the technology with the fourth-highest alert count.

The other four technologies decreased in alert count QoQ. Defined Rules swapped places with Digital Signatures to take the last spot. It dropped the most from Q2, seeing a reduction of 83.47%. Manual Attestation invocations reduced the second-most, decreasing by 24.20%, which our analysts surely appreciate. The other two technologies that had reduced alert counts from Q2 are behavioral/machine learning and Cloud, dwindling by 19.09% and 9.93%, respectively. You can see the spread of these alerts in the Alerts by Technology bar chart (Figure 36).

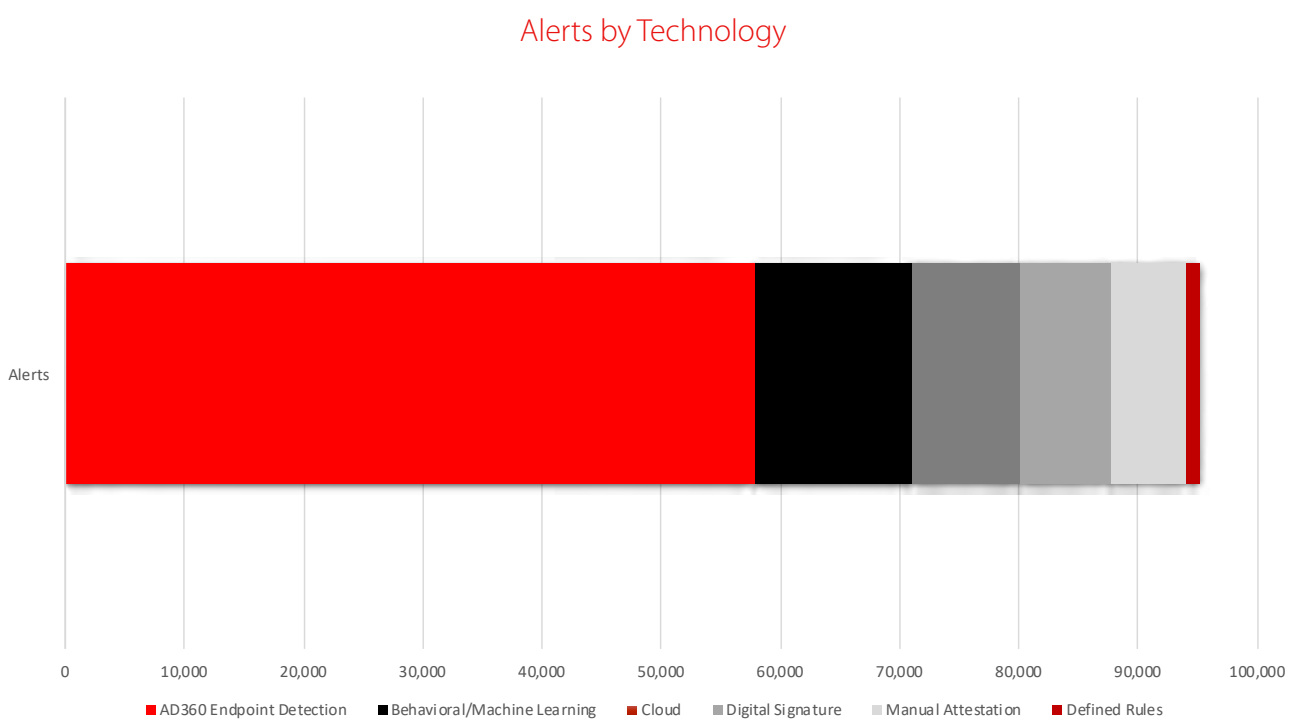

Figure 37. Alerts by Technology

ATTACK VECTORS

Attack Vectors is the longest-tenured subsection of the Endpoint section. It has evolved but maintains the same data points and graphical representation. However, for the first time, we are introducing a new graph that shows the annual timeline of attack vector composition. This new graph allows readers to see all attack vectors simultaneously and which are trending up or down from quarter to quarter. We added it, first, to provide new readers with an understanding of what we mean when we say attack vectors, and for returning readers, a refresher of the attack vector descriptions.

Attack Vector Descriptions

Browsers – Internet browsers are familiar products for all modern-day computer users that allow access to the World Wide Web (WWW). Common browsers include Chrome, Firefox, Safari, and Edge, among many others. Current browsers store personal information – if you allow them – including passwords, cookies, cryptocurrency private keys, and even credit cards. Making them common targets for information-stealing malware.

Office – Office software is the sum of all detections derived from Microsoft Office executables. This includes Word, Excel, PowerPoint, Outlook, and Office Suite executables. Not only is Microsoft Office one of the most popular business-related suites of tools, but the features of the software, such as macro-enablement, allow for an increased attack surface.

Other – The Other attack vector is everything else. Detections within this category are those that did not fit any other category. This includes AutoKMS tools, remote services, and third-party applications, among many others that change every quarter.

Scripts – Scripts, which always invoke the most detections each quarter, are those files derived from or use a scripting programming language. Malware utilizes PowerShell, Python, Bash, and AutoIT scripts to download other malware and deliver payloads, among many other things. Considering Windows is the most commonly attacked operating system, it is no wonder PowerShell continues to skew the results for Windows detections.

Windows – Under the hood, Windows-based software houses the most data points of any attack vector. It contains the most detections but not in the highest quantities. The files included under the Windows name ship with the Windows operating system. Examples include explorer.exe, msiexec.exe, rundll32.exe, and notepad.exe. Trojans commonly impersonate these files or inject malicious code into them because they exist on every Windows machine out of the box and are inherently trusted.

Acrobat has teetered on the edge of inclusion within the described attack vectors for the past few quarters. In fact, there were exactly zero Acrobat-based alerts for the first time last quarter. In Q3, there were some, but, to be frank, it was less than 100. Therefore, we have decided to omit the Acrobat attack vector again and bundle it within the Other attack vector. Because there were so few Acrobat-based alerts, it hardly made any difference to the final numbers.

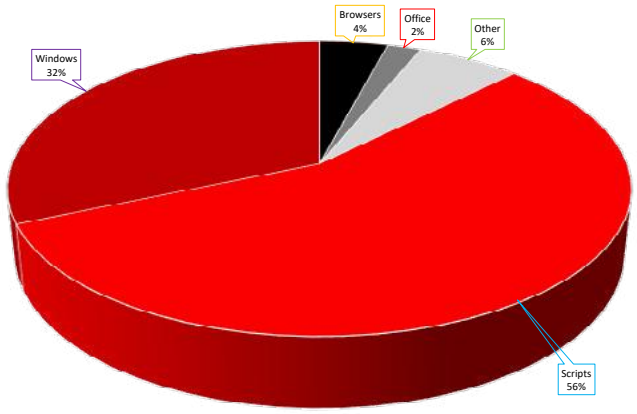

Figure 38. Top Exploited Software

There was only one attack vector that decreased from Q2: Scripts. However, Scripts is by far the attack vector with the most alerts. Surprisingly, Scripts has declined QoQ for several consecutive quarters now. In 2022, the number of Scripts alerts was in six figures. In Q3, the number is now 25,203, comprising only 56% of all alerts. It was farfetched for Scripts to consist of more than 90% of all alerts. So, we've been observing an undeniable decline.

Simply put, we are seeing fewer PowerShell scripts on the endpoint, but it doesn't mean that many of these scripts aren't still alerting on endpoints. Don't be deceived. Remember, Scripts are responsible for over half of all alerts (56%), and an overwhelming majority of those are PowerShell.

Every other attack vector increased between 30% and 60%. The Other attack vector, with contributions from Acrobat, increased by 34.92%, comprising 6% of all alerts. Next is the Browsers attack vector, consisting of 4% of all alerts and rising by 41.74% QoQ. The second-most growing attack vector was Office at 52.82%. However, Office is the smallest data set, with only 2% of all alerts in Q2. Finally, the attack vector that increased the most in Q3 was Windows, and the jump was substantial. Windows increased by 55.62% from Q2 and now is responsible for 32% of all alerts. Last quarter, Windows consisted of about 20% of all alerts. You can observe that significant increase in our new graph: 2023 QoQ Attack Vectors Percentage Totals (Figure 39).

Attack Vector	Q2 Count	Q3 Count	Raw Difference From Q2	Percentage Difference From Q2
Browsers	1093	1876	783	41.74%
Office	435	922	487	52.82%
Other	1851	2844	993	34.92%
Scripts	28046	25203	-2,843	-11.28%
Windows	6280	14152	7,872	55.62%

Figure 39. Attack Vectors

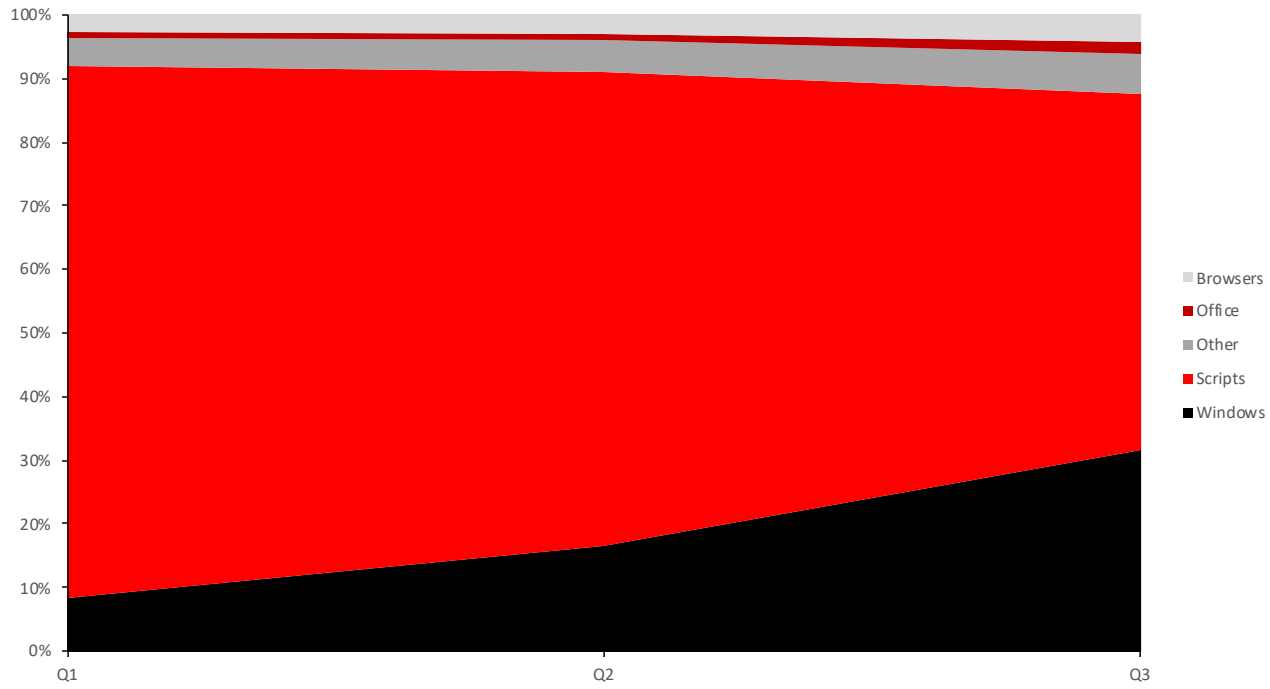

Figure 40. 2023 QoQ Attack Vectors Percentage Totals

Browser Attack Vectors

Another subsection we often couple with the Attack Vectors subsection is the Browser Attack Vectors dissection. We filter all the browser-based detections by which browser family threw the alert. Sometimes, we get browsers many have never heard of or used. For example, we've seen alerts from Opera, Brave, and Edge in the past. However, this quarter, there are only three – Chrome, Internet Explorer, and Firefox. Chrome leads the way with 56% of all browser alerts, followed by Internet Explorer at 33% and Firefox at 11%. This data makes sense, as Chrome is today's most widely used browser. However, we found it odd that Internet Explorer had 33% of alerts and zero Edge alerts.

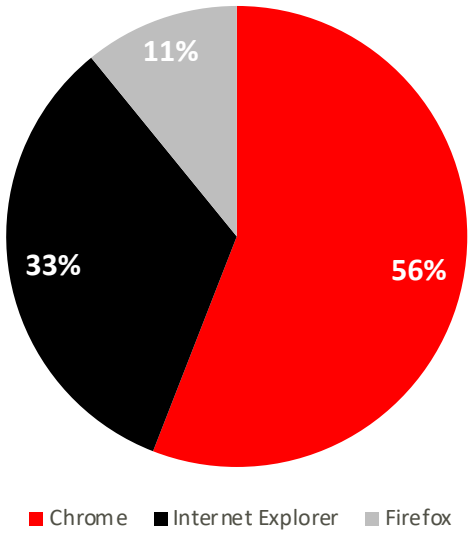

Figure 41. Comparative Browser Detections

Alerts by Exploit Type

The Alerts by Exploit Type subsection dives deeper than the prior subsection. Attack Vectors describes which software types attackers leverage to arrive, pivot, and spread on an endpoint. Exploit types attempt to provide insight into what techniques these attackers use to exploit this software. This list is static for the most part, but occasionally, there is a new exploit type or two that makes the cut. For example, this quarter, there is one new exploit type – AmsiBypass. AMSI stands for Antimalware Scan Interface and is the service within Windows that scans files and determines their maliciousness. Software engineers and other coders can interface with AMSI and develop scripts or other services that call the AMSI service, allowing them to detect malware with Windows services early on in the exploit chain.

Aside from the new exploit type, the rest of the exploit rankings are a mixed bag in reference to Q2. The top two exploits stayed the same, with ShellcodeBehavior and NetReflectiveLoader ranked first and second, respectively. DynamicExec and CVE-2021-26411 also stayed the same as last quarter. The rest moved up or down one or two spots besides RemoteAPCInjection and PsReflectiveLoader1. RemoteAPCInjection moved up five spots, and PsReflectiveLoader1 increased by the most in Q3 by nine ranks. You can see each exploit's order difference and descriptions in the Alerts by Exploit Type table (Figure 41).

Exploit	Alert Count	Description of Exploit
ShellcodeBehavior	12,464	.NET files that allocate and inject payloads directly within the memory of it's own process (Assembly.Load)
NetReflectiveLoader	11,670	Code execution on MEM_PRIVATE pages that do not correspond to a PE
RemoteAPCInjection	8,725	Remote code injection via APCs
RunPE	5,326	Process Hollowing Techniques
ThreadHijacking	2,305	A process injection technique that allows the execution of arbitrary code in a separate process
WinlogonInjection	1,907	Remote Code Injection into winlogon.exe process
PsReflectiveLoader1	1,747	Files that leverage PowerShell to allocate and inject payloads directly within the memory of it's own process (E.g. Mimikats) (Local)
ROP1	775	Return Oriented Programming
IE_GodMode	406	GodMode technique in Internet Explorer
DumpLsass	342	LSASS Process Memory Dump
HookBypass	146	Detection of memory allocation in base addresses; typical of heap spraying
DynamicExec	114	Execution of code in pages without execution permissions (32 bits only)
APC_Exec	43	Local code execution via APC
JS2DOT	39	.NET Reflective Loading Technique
AmsiBypass	20	Techniques that bypass Windows' Antimalware Scan Interface (AMSI)
ReverseShell	20	Detection of reverse shell
ReflectiveLoader	15	Reflective executable loading (Metasploit, Cobalt Strike, etc.)
CVE-2021-26411	4	Microsoft Internet Explorer Memory Corruption Vulnerability

Figure 42. Alerts by Exploit Type

THREAT HUNTING

Last quarter, we omitted the Threat Hunting section due to a data discrepancy. We weren't sure whether the discrepancy was in Q1 or Q2, so we waited until this quarter to determine where the error occurred. Thankfully, we have our answer, and unfortunately, the discrepancy was in Q1, which is already published and viewable. Therefore, we want to first apologize for posting this data. It was the first iteration since our revamp, and we could not know if an error occurred. This admission also means that we omitted the correct data last quarter. Therefore, to make up for it, we decided to post all three quarters, side by side, to give you an idea of how the data has changed QoQ. We are also introducing a new graph summarizing the total MITRE tactics for each technique and subtechnique. As a reminder, the short descriptions of the MITRE tactics and techniques are below.

Tactics and Techniques

We have mapped our successful threat-hunting efforts to techniques in the MITRE ATT&CK matrix. If you are unfamiliar with that framework, you may want to follow some of their **Getting Started** resources to better understand our references in this subsection. The table and the corresponding chart below display the number of threat-hunting occurrences mapped to its appropriate ATT&CK tactic, technique, and sub-technique. The table column headers are:

MITRE Tactic – The primary tactic used. (e.g., TA0002 is Execution)

MITRE Technique – The technique used. (e.g., TA1059.001 is Command and Scripting Interpreter and PowerShell)

Tactic :: Technique :: Sub-Technique – The combined tactic, technique, and sub-technique.

Technique Count – The number of occurrences for each technique.

Tactic Sum – The sum of all Technique Counts for a given Tactic.

Now that you know what happened with Threat Hunting last quarter, let's jump into the QoQ data. The proceeding three tables are in chronological order, from Q1 to Q3. You can see the data discrepancy by referencing the first table and the Q2 or Q3 tables. The Q1 table numbers are abnormally low; we didn't know that until we received the Q2 data. For example, the total Tactic Sum for MITRE Tactic TA0001 (Initial Access) was 42. In Q2, that number was 117,011, an increase of around 2,786 times! However, it's important to remember that you should avoid analyzing the Q1 table; it's wrong. Plain and simple.

Retrospection aside, the Q3 data from the Q2 data saw a massive increase for each tactic. The tactics TA0001, TA0002, TA0004, TA0005, TA0007, TA0008, and TA0011 all had roughly triple the threat-hunting rule invocations QoQ. Whereas tactics TA0003, TA0006, TA0040 roughly doubled from Q2 to Q3. If an action matches one of our threat-hunting rules, it doesn't mean it is malicious. It more or less means the action is suspicious and needs additional analysis from a WatchGuard threat-hunting team member. Therefore, just because there is a substantial increase in threat-hunting rule invocations, it doesn't necessarily mean more observed malware. The Malware Frequency data supports this claim.

MITRE Tactic	MITRE Technique	Tactic :: Technique :: Sub-Technique	Technique Count	Tactic Sum
TA0001	TA0001	Initial Access	42	42
TA0002	TA0002	Execution	1,296	9,792
	T1059.001	Execution :: Command and Scripting Interpreter :: PowerShell	8,200	
	T1218.011	Execution :: Signed Binary Proxy Execution :: Rundll32	47	
	T1543.003	Execution :: Create or Modify System Process :: Windows Service	109	
	T1569.002	Execution :: System Services: Service Execution :: Service Execution	140	
TA0003	TA0003	Persistence	2,651	2,695
	T1543.003	Persistence :: Create or Modify System Process :: Windows Service	17	
	T1546.008	Persistence :: Event Triggered Execution :: Accessibility Features	11	
	T1546.012	Persistence :: Event Triggered Execution :: Image File Execution Options Injection	10	
	T1547.001	Persistence :: Boot or Logon Autostart Execution :: Registry Run Keys / Startup Folder	6	
TA0005	TA0005	Defense Evasion	344	383
	T1070.004	Defense Evasion :: Indicator Removal :: File Deletion	8	
	T1218.009	Defense Evasion :: System Binary Proxy Execution :: Regsvcs/Regasm	5	
	T1218.011	Defense Evasion :: System Binary Proxy Execution :: Rundll32	6	
	T1562.001	Defense Evasion :: Impair Defenses :: Disable or Modify Tools	20	
TA0006	TA0006	Credential Access	434	636
	T1555.003	Credential Access :: Credentials from Password Stores :: Credentials from Web Browsers	202	
TA0007	TA0007	Discovery	24	24
TA0008	TA0008	Lateral Movement	498	1,070
	T1021.001	Lateral Movement :: Remote Services :: Remote Desktop Protocol	572	
TA0010	TA0010	Exfiltration	6	6
TA0011	TA0011	Command and Control	113	114
TA0040	TA0040	Impact	87	107
	T1561.001	Impact :: Disk Wipe :: Disk Content Wipe	20	

Figure 43. Exploits by MITRE ATT&CK Tactic and Technique Table, Q1 2023

MITRE Tactic	MITRE Technique	Tactic :: Technique :: Sub-Technique	Technique Count	Tactic Sum
TA0001	TA0001	Initial Access	117,011	117,011
TA0002	TA0002	Execution	597,576	3,086,227
	T1059.001	Execution :: Command and Scripting Interpreter :: PowerShell	2,116,689	
	T1543.003	Execution :: Create or Modify System Process :: Windows Service	317,188	
	T1569.002	Execution :: System Services :: Service Execution	54,774	
TA0003	TA0003	Persistence	599,322	1,850,999
	T1053.005	Persistence :: Scheduled Task/Job :: Scheduled Task	714,597	
	T1543.003	Persistence :: Event Triggered Execution :: Accessibility Features	343,136	
	T1547.001	"Persistence :: Event Triggered Execution :: Image File Execution Options Injection	172,868	
	T1547.006	Persistence :: Boot or Logon Autostart Execution :: Kernel Modules and Extensions	21,076	
TA0004	TA0004	Privilege Escalation	26,707	95,087
	T1548.003	Privilege Escalation :: Abuse Elevation Control Mechanism :: Sudo and Sudo Caching	68,380	
TA0005	TA0005	Defense Evasion	1,846,039	2,739,759
	T1027.004	Defense Evasion :: Obfuscated Files or Information :: Compile After Delivery	27,618	
	T1070.004	Defense Evasion :: System Binary Proxy Execution :: Regsvcs/Regasm	482,480	
	T1218.009	Defense Evasion :: System Binary Proxy Execution :: Rundll32	354,101	
	T1562.001	Defense Evasion :: Impair Defenses :: Disable or Modify Tools	29,521	
TA0006	T1552.001	Credential Access :: Unsecured Credentials :: Credentials In Registry	322,450	373,061
	T1552.002	Credential Access :: Unsecured Credentials :: Credentials In Files	33,851	
	T1558.003	Credential Access :: Steal or Forge Kerberos Tickets :: Kerberoasting	16,760	
TA0007	TA0007	Discovery	2,784,853	2,784,853
TA0008	T1021.001	Lateral Movement :: Remote Desktop Protocol :: Remote Services	126,143	126,143
TA0011	TA0011	Command and Control	791,091	791,091
TA0040	TA0040	Impact	159,768	2,363,820
	T1561.001	Impact :: Disk Wipe :: Disk Content Wipe	2,204,052	

Figure 44. Exploits by MITRE ATT&CK Tactic and Technique Table, Q2 2023

MITRE Tactic	MITRE Technique	Tactic :: Technique :: Sub-Technique	Technique Count	Tactic Sum
TA0001	TA0001	Initial Access	385,591	385,591
TA0002	TA0002	Execution	1,580,661	9,174,714
	T1059.001	Execution :: Command and Scripting Interpreter :: PowerShell	6,916,862	
	T1543.003	Execution :: Create or Modify System Process :: Windows Service	559,465	
	T1569.002	Execution :: System Services :: Service Execution	117,726	
TA0003	TA0003	Persistence	1,695,347	4,177,682
	T1053.005	Persistence :: Scheduled Task/Job :: Scheduled Task	1,719,027	
	T1543.003	Persistence :: Event Triggered Execution :: Accessibility Features	292,230	
	T1547.001	"Persistence :: Event Triggered Execution ::	419,753	
	T1547.006	Persistence :: Boot or Logon Autostart Execution :: Kernel Modules and Extensions	51,325	
TA0004	TA0004	Privilege Escalation	51,638	300,511
	T1548.003	Privilege Escalation :: Abuse Elevation Control Mechanism :: Sudo and Sudo Caching	248,873	
TA0005	TA0005	Defense Evasion	4,598,252	6,756,760
	T1027.004	Defense Evasion :: Obfuscated Files or Information :: Compile After Delivery	65,629	
	T1070.004	Defense Evasion :: System Binary Proxy Execution :: Regsvcs/Regasm	1,174,568	
	T1218.009	Defense Evasion :: System Binary Proxy Execution :: Rundll32	844,383	
	T1562.001	Defense Evasion :: Impair Defenses :: Disable or Modify Tools	73,928	
TA0006	T1552.001	Credential Access :: Unsecured Credentials :: Credentials In Registry	677,523	796,848
	T1552.002	Credential Access :: Unsecured Credentials :: Credentials In Files	80,576	
	T1558.003	Credential Access :: Steal or Forge Kerberos Tickets :: Kerberoasting	38,749	
TA0007	TA0007	Discovery	7,763,523	7,763,523
TA0008	T1021.001	Lateral Movement :: Remote Desktop Protocol :: Remote Services	329,982	329,982
TA0011	TA0011	Command and Control	2,021,623	2,021,623
TA0040	TA0040	Impact	374,832	5,693,511
	T1561.001	Impact :: Disk Wipe :: Disk Content Wipe	5,318,679	

Figure 45. Exploits by MITRE ATT&CK Tactic and Technique Table, Q3 2023

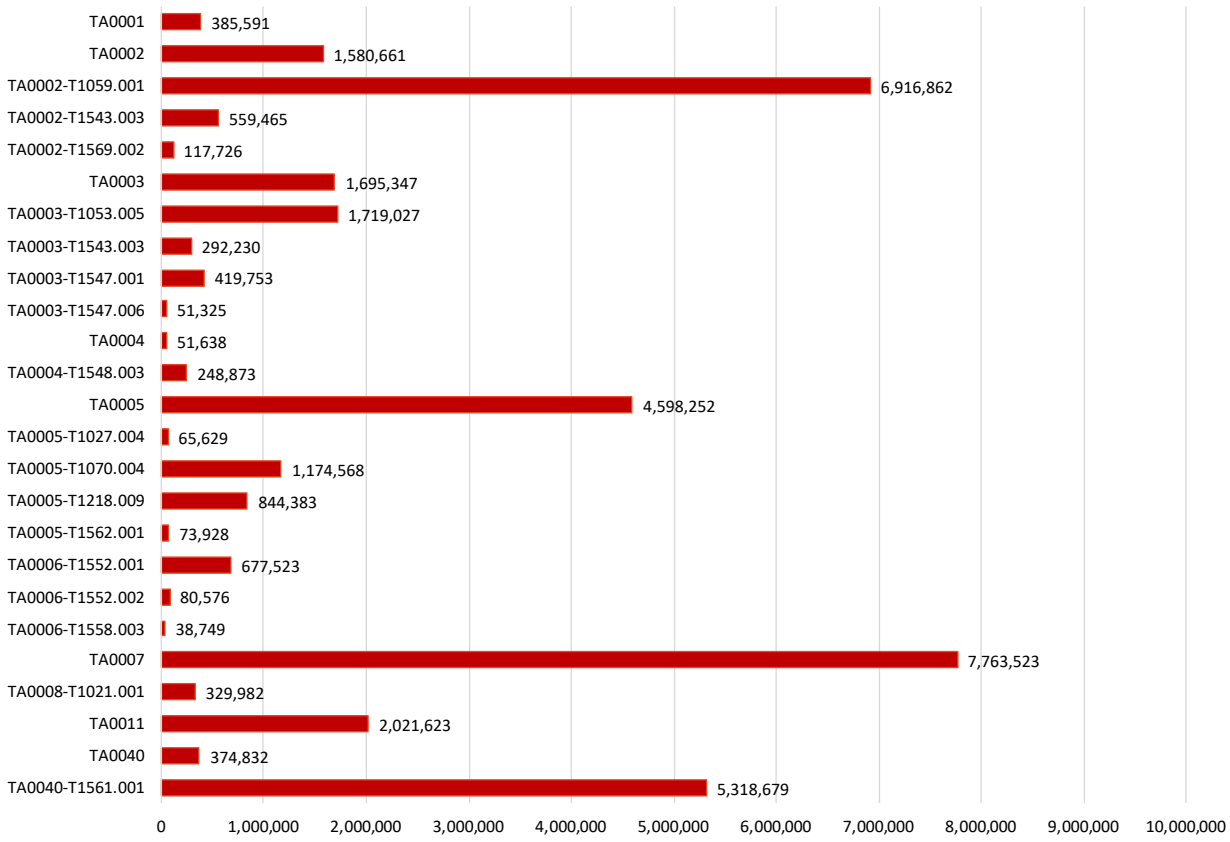

Figure 46. Exploits by MITRE ATT&CK® Tactic and Technique

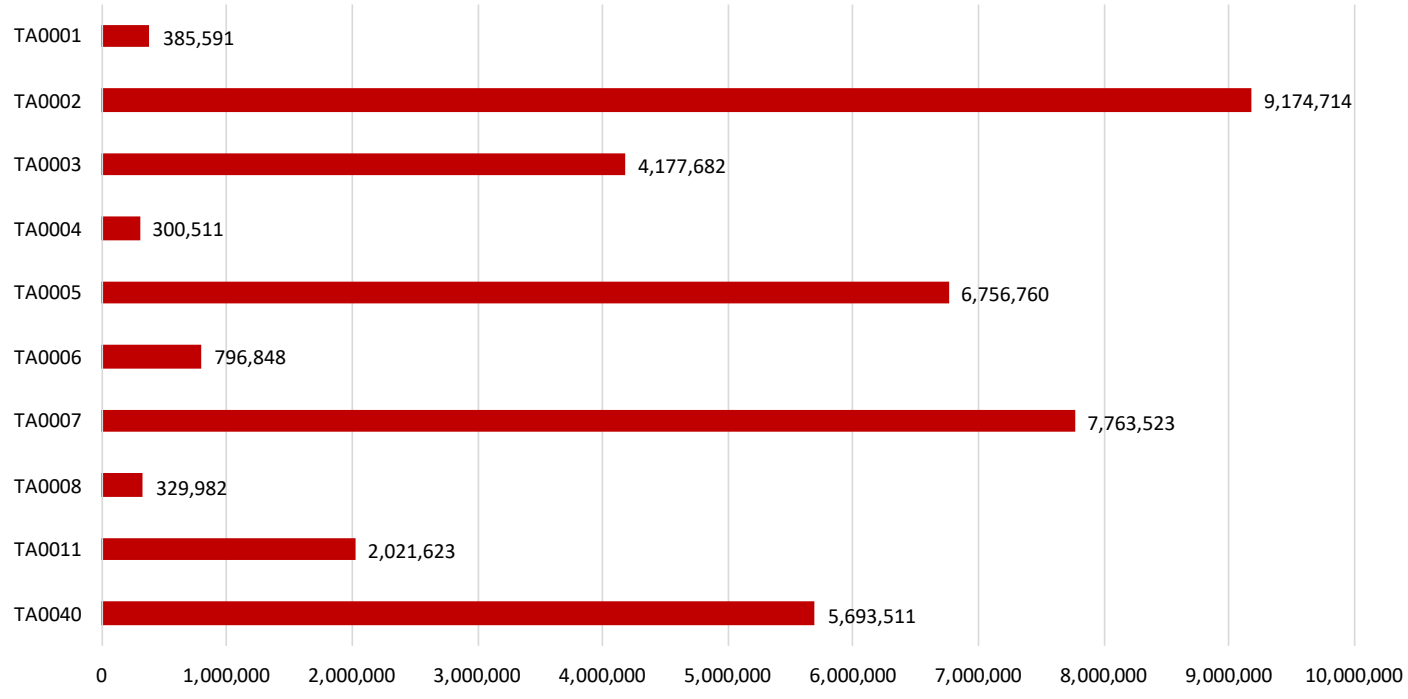

Figure 47. Exploits by MITRE ATT&CK® Tactics Summation

RANSOMWARE LANDSCAPE

The ransomware landscape has been good for the bad guys and bad for the good guys. However, given the ransomware detections on WatchGuard Fireboxes in Q3, you wouldn't believe that. After a significant increase in ransomware detections in 2022, 2023 has shown a QoQ decline, which continues into Q3. We only observed 421 ransomware detections in Q3, a slight decrease of 9.46% from Q2. This decline follows the trend from Q1 to Q2, which decreased 21.58%. Overall, ransomware detections on the Firebox are down 29% for the year, but we believe this number has likely bottomed out. Our theory is that the detections are so low because ransomware operators perform various other malicious actions before deploying a ransomware payload. Thus, there's a good chance of being caught before the payload can reach the endpoint.

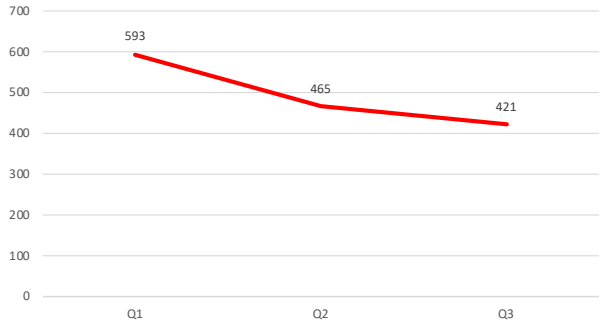

Figure 48. Ransomware Detections by Quarter

Extortion Groups

The earlier claim that the ransomware landscape is going well for the bad guys is due to the double extortions observed in Q3. This number provides a more all-encompassing analysis of ransomware operator's fairing this quarter. Double extortion is a tactic that began with the Snatch group a few years ago and has now become the norm for ransomware groups. Once they gain access to a network, they attempt to exfiltrate as much data as possible and sometimes will begin to negotiate for the return or deletion of this data. If the victim doesn't cooperate or pay the ransom demand, the group will often post the data on a data leak site to shame them into paying (double extortion). If they still don't pay, the data is published publicly or sold to the highest bidder on dark web forums.

In Q3, several new groups dashed onto the scene:

New Groups:

- Cactus
- CiphBit
- Cloak
- CryptBB
- Cyclops/Knight
- DataLeakes
- INC Ransom
- LostTrust
- Metaencryptor
- NoEscape
- RansomedVC
- ThreeAM

There are a few interesting things to note from these new groups. First, Cyclops appeared in Q3, but shortly after their discovery, they rebranded to Knight. This is why you see the group shown as Cyclops/Knight. It's two different names for the same group. Second, it appears that LostTrust and Metaencryptor are by the same operators. Their dark web data leak sites look the same. Albeit, the posted extortions are different. Surprisingly, 8base and the new CryptBB group follow the same pattern; they appear to be run by the same operators. Finally, the last notable mention from the new groups is the NoEscape operation. Researchers believe this is a rebrand of the infamous Avaddon group from years prior.

Ransomware detections for EPDR users are down for Q3, but so are double extortions, but ever so slightly. We observed a 6.23% decrease in double extortion attempts by ransomware operators from Q2 to Q3. Even though the numbers are down from Q2, the overall numbers for Q3 are still elevated far above Q1 levels. Double extortions are still up 54.35% from Q1, even down from Q2. There was a decrease, but some groups performed more extortions in Q3 than in Q2. The groups that had an increase of victims from the quarter prior are:

New Groups:

- Abyss
- Arvin Club
- CL0P Leaks
- CryptBB*
- Cyclops/Knight*
- DataLeakes*
- DungHill Leak
- Everest
- LockBit 3.0
- MedusaLocker
- NoEscape*
- Play
- RA Group
- Ragnar Locker
- Rhysida
- Stormous

The groups with an asterisk are those we learned about in Q3 but had extortions dating back to Q2 or earlier. So, they are both new and had historically fewer extortions than they did in Q3. To fortify understanding of this topic, we have introduced a new graph that shows the annual double extortions for each group while simultaneously showing how each quarter's numbers contributed. You can view that below in the red Public Extortions by Group graph.

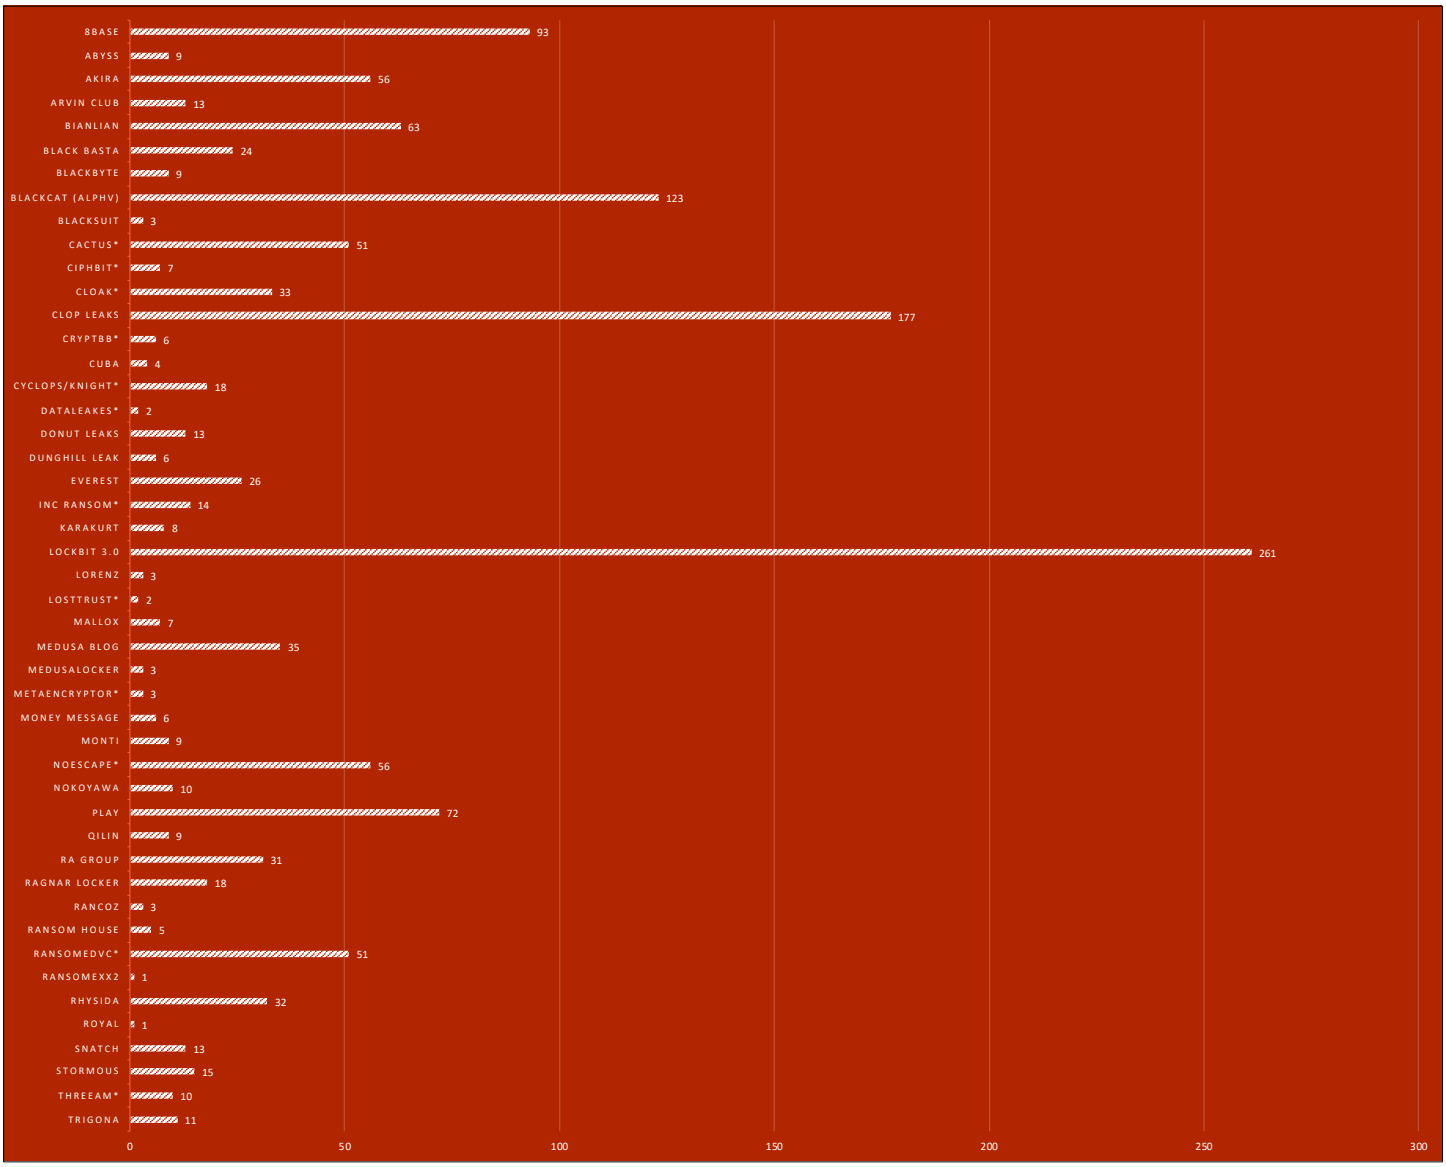

Figure 49. Public Extortions by Group

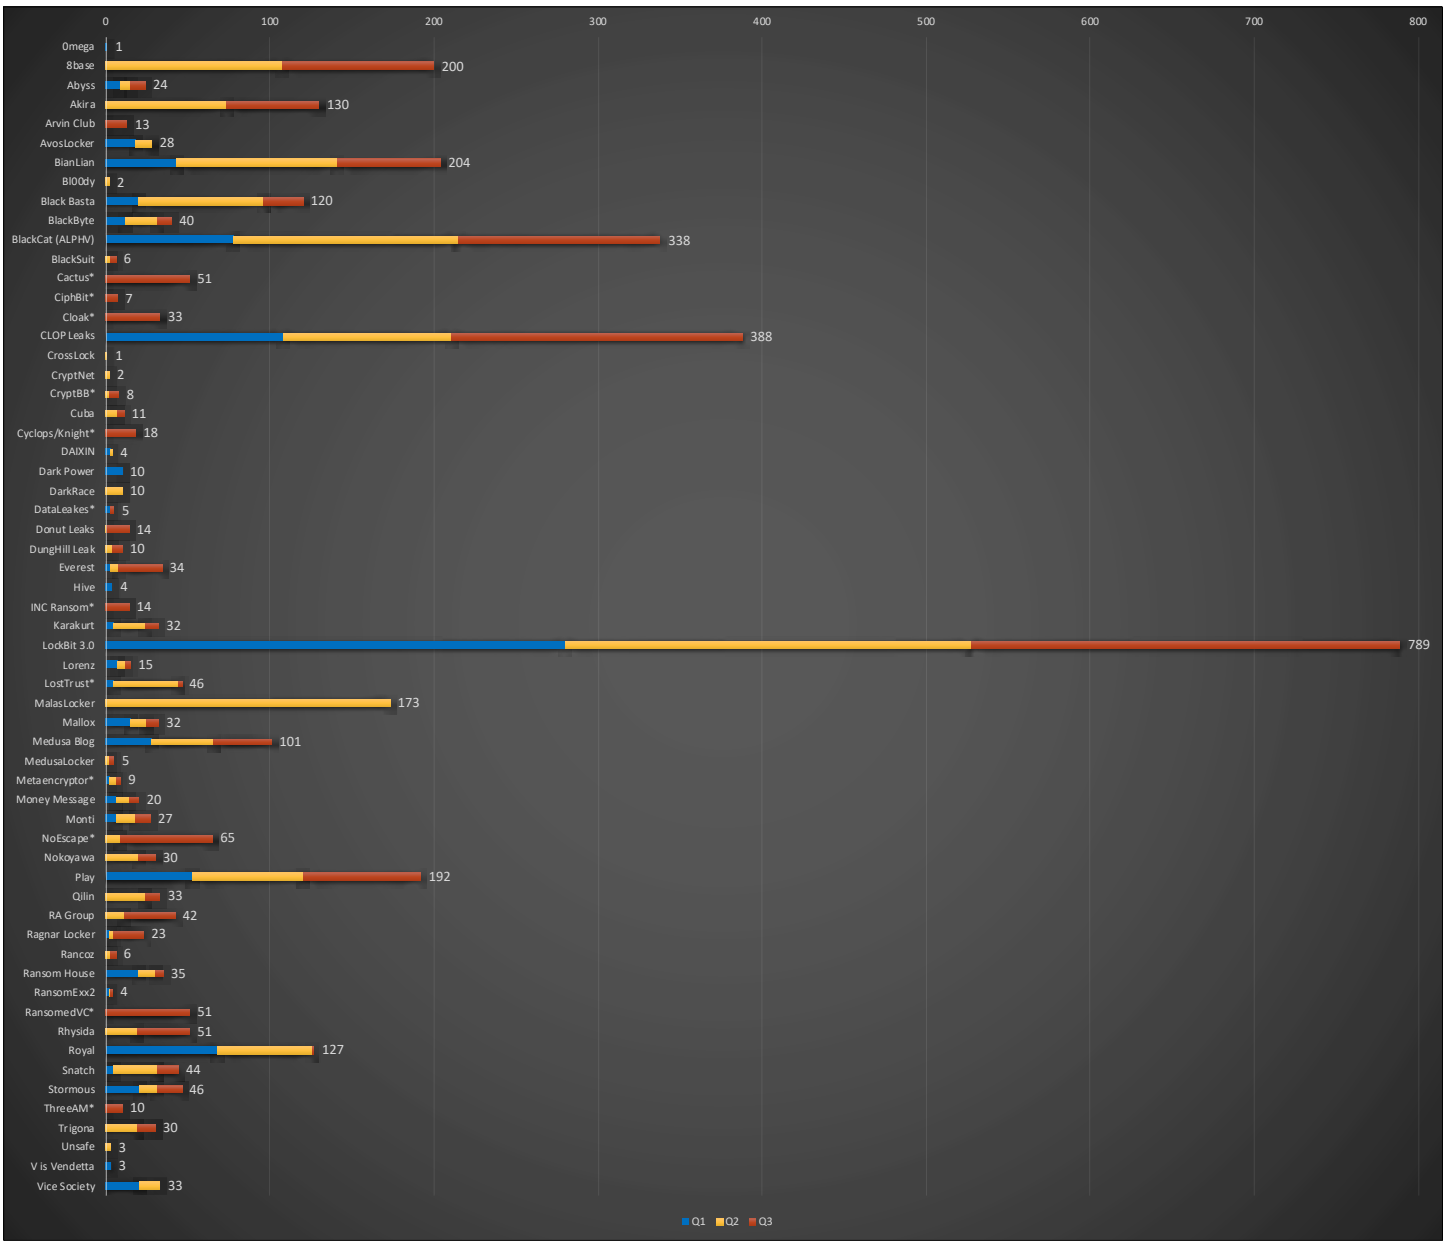

Figure 50. 2023 QoQ Public Extortions by Group Summation

Notable Ransomware Breaches

Akira

KNP Logistics – Unfortunately, the first notable breach for Q3 is more tragic than a simple ransomware double extortion scheme. Towards the end of September, details emerged of a ransomware attack on KNP Logistics. The Akira group listed this organization in early August. However, it wasn't until late September that we learned most of the details of this attack. At first, this was yet another ransomware attack on a major distribution and logistics company. KNP Logistics is one of the largest private logistics companies in the United Kingdom. This fact alone makes this breach notable. However, decision-makers sold most of the company, leaving only around 15% of the workforce. Executives from the company blamed the ransomware attack as the primary reason for the significant organizational shift. This breach proves that ransomware attacks can accomplish much more than a financial hit to organizations.

BianLian

Save the Children International – This breach is listed because of who it is and not so much as what happened. Save the Children International is a nonprofit organization focusing on children, as the name implies. The organization began over 100 years ago, right after World War 1, by a woman named Eglantyne Jebb. Her goal was simple: to ensure children don't experience the hardships experienced during the war. Today, they operate globally and assist children in all circumstances – war, famine, natural disasters, disease, and much more. BianLian breached this charity and posted them on their dark web data leak site. Although they didn't explicitly list them, the group attempted to mask the extortion as **** **e ******e* ***e********* (Save The Children International). In the end, BianLian claimed to have 6.8 TeraBytes (TB) of data, but the charity said the attack didn't impact any operations.

BlackBasta

BankCard USA (BUSA) – As the name implies, BUSA is an American-based company that provides end-to-end payment services to thousands of other American companies. Unfortunately for BUSA, researchers published much of the details of this breach, including BUSA allegedly paying a $50,000 ransom to Black Basta in exchange for the group to destroy the stolen data. However, Black Basta posted the group to their dark web data leak site with some of the alleged stolen data. In other words, they lied, a prime example of not trusting these ransomware extortion groups.

BlackCat (ALPHV)

MGM Resorts and Caesars Entertainment – If you followed cybersecurity-related news in Q3, you've undoubtedly heard of this duo of breaches. At the beginning of September, reports began to flood the news of a widespread cyberattack occurring in Las Vegas. These reports and videos from social media users showing various slot machines and other gambling devices out of order confirmed these findings. At this point, it was undeniable that something serious was happening, and very few cyberattacks cause widespread outages such as these – the most apparent being ransomware.

It didn't take long to learn that this was a ransomware attack, and the ALPHV group claimed responsibility. However, Caesars Entertainment avoided much of the destruction as they allegedly paid a ransom before an encryption event. Therefore, Caesars didn't have a hindrance in their operations as much as companies under the MGM Resorts umbrella. MGM Resorts did experience an encryption event from the ALPHV group, which supports the evidence of gambling machines being out of order. Astonishingly, researchers claim that an ALPHV affiliate dubbed Scattered Spider breached these two organizations by using social engineering attacks via LinkedIn, and the attack took mere minutes to execute. This is yet another example of how easy it can be to breach a network, no matter what technological solutions you have in place.

Dark Angels

Johnson Controls International – At the very end of Q3 (September), reports began to swirl that Johnson Controls International, a global conglomerate that manufactures industrial control systems and other equipment for organizations, became a victim of a cyberattack. We quickly learned that this was a ransomware attack and the threat actors encrypted several of the company's servers, including VMware ESXi servers. Based on the ransom note used in the attack, which researchers managed to get a hold of, the Dark Angels group was responsible for the attack. This group also manages another operation, Dunghill Leak, which claimed several victims in Q3. However, the Dark Angels group used another self-named data leak site – Dark Angels – for this breach. We listed Johnson Controls International as a notable breach for Q3 because of the widespread reach as a manufacturer, their employment of over 100,000 people, and because reports claim that the ransomware operators demanded $51 million in payment.

Karakurt

McAlester Regional Health Center – This alleged breach by Karakurt is an example of a worst-case scenario for patients. The Karakurt group posted McAlester Regional Health Center on August 1, 2023, and claimed to have exfiltrated 126 gigabytes (GB) of data from the health center, including 40 GB of patient DNA tests. Considering that these groups tend to sell this information to the highest bidder if the victim doesn't send a ransom payment, patients could see their DNA information sold to other nefarious threat actors, including entities tied to nation-states. It's unclear if McAlester decision-makers paid Karakurt or if the group successfully sold this information, but it is clear that the patient's most sensitive information from this health center is possibly on the Internet somewhere.

NoEscape

Au Domain Administration (auDA) – The auDA is the custodian of the Australian domain namespace, as their company name implies. Any attack on a company that performs essential tasks for the availability of Internet domain resolution has severe consequences for users in Australia. At first, the auDA denied any cyberattack on their network. However, when the NoEscape ransomware group posted proof of the attack, the auDA admitted it. According to the NoEscape group, the negotiations between the two entities went sour. Instead of giving the victim eight days for payment, they lowered it to three and, after more mishaps, reduced it to two. Additionally, the group said they would begin auctioning off bank accounts with more than a $4k balance. As a silver lining, there were no reports of operational failure or availability lapses.

Rhysida

Prospect Medical Holdings – Prospects Medical Holdings is a major healthcare provider with hundreds of medical and outpatient centers in the United States. At the beginning of August, employees of the medical conglomerate reported seeing a ransom note on their endpoints when arriving at work. The note didn't explicitly state who was behind the attack, but the dark web domain in the ransom note gave it away – it was the Rhysida ransomware group. The group claims to have exfiltrated hundreds of thousands of corporate documents, including personally identifiable information and social security numbers. Additionally, the attack caused so much damage that employees had to record patient information with pen and paper until the IT group restored systems to normal. It's not the worst-case scenario, but it's close to it.

Kuwait Ministry of Finance – The Rhysida group is making a name for itself, having gone after government agencies in several countries, including South America, Asia, and the Caribbean. This victim is yet another example. On September 25, The official Twitter (X) account of Kuwait's Ministry of Finance disclosed they were subject to a ransomware attack. The statement also mentioned they were proactively isolating systems and bringing external assistance to alleviate the problem. The same day, the Rhysida operators listed the ministry on their dark web data leak site, giving them seven days to pay an unknown financial demand. It is unclear if they ever paid the extortion.

Snatch

South Africa Department of Defence (DARPA) – Snatch has an extensive history in the ransomware double extortion space. They were the first to post a double extortion on the dark web. Unfortunately, they are still operating and extorting victims. Even more unfortunate, the victim listed here is a notable breach in Q3 because it's a significant entity in South Africa. Having ransomware deployed on the Department of Defence of a country is a scenario many citizens would not favor because of the obvious – it's the department in charge of the defense of a nation. A significant and widespread attack could cripple military and utility infrastructure. The Snatch group claims to have stolen 1.6 terabytes (TB) of data from the department, including personal information and defense contracts. If a particular contract gets into the wrong hands, it could spell more trouble. Fortunately, there was no immediate impact on any infrastructure of the department.

Unknown Group

CloudNordic and AzeroCloud – CloudNordic and AzeroCloud are sister companies offering similar solutions to users in Denmark. They are both Cloud-hosting providers and host almost all their user's data. Knowing this tidbit of information, a cyberattack on either of these institutions could spell disaster for users. Discouragingly, that's precisely what happened. Not only that, but both organizations had their backups and secondary backups encrypted, too. Therefore, they couldn't "restore to backups" and lost almost all their data, including their customers. Instead of paying the ransom demand, they took the moral high ground, refused to pay the cybercriminal's demand, and subsequently shut down their operations. Thus, these two companies are no more, and which group was responsible for these attacks is unknown.

CONCLUSION & DEFENSE HIGHLIGHTS

CONCLUSION AND DEFENSE HIGHLIGHTS

That summarizes the threat landscape patterns and findings we discovered in Q3. To summarize, malware is up, especially evasive zero-day. Network attacks are also up, with headline-grabbing exploits like ProxyLogin topping our lists. Threat actors seem to still prefer living-off-the-land (LotL) attack techniques as both scripts and commonly exploited Window binaries are the two most regular vectors for endpoint-based malware. Finally, cybercriminals still like targeting open remote access products, or like to leverage legitimate remote access tools to hide their malicious actions. But don't worry, there are solutions for these threats.

We already gave you some of those tips throughout this report, but let's end with three final strategies that can protect you from the most common attack patterns of Q3 2023.

Locking down remote access aggressively

According to the top security pros, government agencies, and even insurance companies, remote access software is one of the top risks organizations face. Whether it's remote desktop protocol (RDP), virtual private networks (VPN), remote monitoring and management (RMM) software, or one of the many screen-sharing apps out there like VNC, TeamViewer, AnyDesk, GoToMyPC and countless more, threat actors have breached many networks via exposed remote access apps and lost, stolen, or cracked credentials. Even if you haven't exposed a remote access app yourself, many social networking attacks try to trick your users into installing a perfectly legitimate one, but with configurations that give them access. For instance, during Q3 we saw multiple phishing domains spread a legitimate version of TeamViewer with a configuration file that gave a criminal access to the machine of anyone who installed it.

This is why it's critically important for you to lock down all remote access apps! So how do you do that? Here are a few tips:

- ***Do not expose your intended remote access apps to the Internet without considerations and protections.*** In general, you should not expose RMM, management interfaces, or any remote desktop app to all people on the Internet. Instead, only allow users to access it through VPN to offer additional protection and security. Sure, you will probably have to allow VPN access globally (though you could also limit it to an access list too), but it tends to be better hardened for Internet exposure, and we highly recommend you only allow VPN with multi-factor authentication (MFA) enabled.
- ***Scan for accidental remote access exposure.*** Even if you do need to expose some remote access (like VPN), you should know all the remote access apps exposes. To verify that is all you've exposed, you can use vulnerability assessment or port scanning software to scan your network and identify any open remote access services. Many endpoint products have application detection and control services that can tell you if computers have remote access apps installed and listening on the network. If you find any remote access services you do not expect, get rid of them.
- ***Leverage application white- or blacklisting.*** Even if you decide to allow some remote access products, you should standardize on which tools and only allow those to run. Endpoint protection suites, like WatchGuard EPDR, often allow you to blacklist remote access apps you aren't using, even if they are legitimate. For instance, if your IT organization standardizes on TeamViewer for some remote access, that is the only remote access product you should allow. Blacklist all the rest.

Patch Microsoft products diligently

We often remind you that patching software is one of the best security strategies, since most network attacks leverage old software that already has a patch. Attackers can't exploit the bugs that you have fixed, and most of the network attacks we have seen in our report received fixes long ago. So, apply those patches.

However, this advice is especially important with Microsoft products, as they are among the most targeted ones by threat actors. In Q3, the top attempted exploit was against the ProxyLogin vulnerably affecting Microsoft Exchange servers. We also saw criminals exploiting ProxyNotShell. Both those issues were fixed long ago, as long as you downloaded and installed the updates. Microsoft Patch Day occurs without fail every second Tuesday of the month. Be sure to watch for that day and apply its security updates as quickly as you can.

Scan and/or strip dangerous email attachments

In this report, you learned about a malware dropper, Stacked, that had many variants on both our Top 10 malware and Top 5 widespread malware lists. This threat arrived as an attachment in emails. Different Stacked variants may arrive as a .RAR file, an .EXE, a document, or PDF, often masquerading as an invoice file; but however it arrives you can implement some basic email security to prevent these malicious files from getting through.

If you have your own email server, you can use the Firebox SMTP proxy to scan all email attachments with our many anti-malware services, thus catching and removing the malicious Stacked files. You can also use our extension-blocking features to completely remove all unnecessary attachments from email, like EXE files, .HTM files, and .LNK files. Why allow any of these files that have no legitimate reason to arrive in email?

If you are using Microsoft M365 Cloud email, you can even use WatchGuard's Email Protection service to also scan for and remove malicious attachments. In short, much of the top malware we see still arrives as simple attachments in email, so be sure to completely strip attachments you don't want to see in email, and leverage email malware scanning services to catch the rest.

We hope you found our Q3 2023 Internet threat landscape report interesting, and maybe gleaned a new tip or two. Return next quarter to see how the patterns and trends continue or change. As always, leave your comments or feedback about our report at **SecurityReport@watchguard.com**, and keep frosty online!

COREY NACHREINER

Chief Security Officer

Recognized as a thought leader in IT security, Corey spearheads WatchGuard's security vision. Corey has operated at the frontline of cybersecurity for 22 years, evaluating and making accurate predictions about information security trends. Corey has the expertise to dissect complex security topics, making him a sought-after speaker at forums such as Gartner, Infosec and RSA. He is also a regular contributor to leading publications including CNET, Dark Reading, Forbes, Help Net Security, and more. Find him on www.secplicity.org.

MARC LALIBERTE

Director of Security Operations

Specializing in network security technologies, Marc's industry experience allows him to conduct meaningful information security research and educate audiences on the latest cybersecurity trends and best practices. With speaking appearances at IT conferences and regular contributions to online IT and security publications, Marc is a security expert who enjoys providing unique insights and guidance to all levels of IT personnel.

TREVOR COLLINS

Information Security Analyst

Trevor Collins is a information security analyst at WatchGuard Technologies, specializing in network and wireless security. Trevor earned his security know-how and several certifications through his past military experience in the United States Air Force. Trevor is a regular contributor to Secplicity.org where he provides easily understood data analysis and commentary to IT professionals. Trevor's experience with a wide range of network security vendors and technologies allows him to provide unique perspectives to the industry.

RYAN ESTES

Intrusion Analyst

Ryan is an intrusion analyst at WatchGuard Technologies operating primarily within DNSWatch, WatchGuard's DNS filtering and security service. For DNSWatch, Ryan helps customers better understand potential threats to their organization using tailored domain analysis and threat intelligence. Outside of DNSWatch, his research interests include web application security, Wi-Fi communications, and malware analysis. Ryan embraces a 'never stop learning' lifestyle allowing him to stay on top of the latest cybersecurity and malware trends. In turn, Ryan passes this knowledge on to our customers and even shares certain topics on his personal blog.

JOSH STUIFBERGEN

Intrusion Analyst

Josh is an intrusion analyst at WatchGuard Technologies operating primarily within DNSWatch, WatchGuard's DNS filtering and security service. For DNSWatch, Josh helps customers better understand potential threats to their organization using tailored domain analysis and threat intelligence. Josh's multidisciplinary background with a political science BA and cybersecurity BS offers an added perspective into the geopolitical nature of cybersecurity threats. Past experience researching container security in Kubernetes deployments, and building a Zero-Trust Proof of Concept environment for small organizations contributes to his insights on how organizations face the difficulties of increasingly complex managed environments. His role includes contributing to the Secplicity blog.

ABOUT WATCHGUARD THREAT LAB

WatchGuard's Threat Lab is a group of dedicated threat researchers committed to discovering and studying the latest malware and Internet attacks. The Threat Lab team analyzes data from WatchGuard's Firebox Feed, internal and partner threat intelligence, and a research honeynet, to provide insightful analysis about the top threats on the Internet. Their smart, practical security advice will enable you to better protect your organization in the ever-changing threat landscape.

ABOUT WATCHGUARD TECHNOLOGIES

WatchGuard® Technologies, Inc. is a global leader in unified cybersecurity. Our Unified Security Platform® approach is uniquely designed for managed service providers to deliver world-class security that increases their business scale and velocity while also improving operational efficiency. Trusted by more than 17,000 security resellers and service providers to protect more than 250,000 customers, the company's award-winning products and services span network security and intelligence, advanced endpoint protection, multi-factor authentication, and secure Wi-Fi. Together, they offer five critical elements of a security platform: comprehensive security, shared knowledge, clarity & control, operational alignment, and automation. The company is headquartered in Seattle, Washington, with offices throughout North America, Europe, Asia Pacific, and Latin America. To learn more, visit WatchGuard.com.

For additional information, promotions and updates, follow WatchGuard on Twitter @WatchGuard, on Facebook, and on the LinkedIn Company page. Also, visit our InfoSec blog, Secplicity, for real-time information about the latest threats and how to cope with them at www.secplicity.org.