

INTERNET SECURITY REPORT

Q2 2023





CONTENTS

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

- 03 Introduction
- 04 Executive Summary
- 06 Firebox Feed Statistics
 - 08 Malware Trends
 - 09 Top 10 GAV Malware Detections
 - 10 Top 5 Encrypted Malware Detections
 - 10 Top 5 Most-Widespread Malware Detections
 - 11 Geographic Threats by Region
 - 11 Catching Evasive Malware
 - 12 Individual Malware Sample Analysis
 - 14 Network Attack Trends
 - 15 Top 10 Network Attacks Review
 - 18 Most-Widespread Network Attacks
 - 20 Network Attack Conclusion
 - 21 DNS Analysis
 - 21 Top Malware Domains
 - 23 Firebox Feed: Defense Learnings
- 24 Endpoint Threat Trends
 - 27 Top Malware and PUPs
 - 31 Attack Vectors
 - 33 Ransomware Landscape
- 36 Conclusion and Defense Highlights
- 39 About WatchGuard

INTRODUCTION

"I like stirring the pot—I think it's part of my duty to shake people up a bit—make them look at things in a different way."

~ Nina Bawden (unverified)

We stirred the pot.

Last quarter, we greatly changed some of the methodologies we use to normalize our network malware data to hopefully discard less meaningful outliers to get a better view of the actual cybersecurity trends you can learn from. This quarter, we continue stirring the pot with the expansion of those methodologies to our network attack and endpoint malware sections too.

Our intention in stirring our data is to find the most accurate results and perhaps uncover new trends and learnings from our view of the threat landscape. However, when you first stir a pot, you agitate its ingredients and make the soup cloudier. That is where we are with these new results. They have helped us see our security trends in a different way, but they also reset our historical view, making it a bit harder for us to interpret results that differ from the past. As with a stirred pot of stew, we expect the cloudy broth to eventually clear up as trends settle to this new normal. But for now, bear with us as we build our understanding and a new history with these new results.

Stirred pot or not, the resulting stew is still good. Like every quarter, we gather and aggregate some of the important threat intelligence we get from various WatchGuard network and endpoint products to identify the threat landscape trends you should know about so that you can defend against them. We look at prominent malware, the most attacked network services, endpoint security trends, and more, hoping to give you some idea of what cybercriminals have been doing, and our estimates of how those trends might evolve in the future. Us stirring the pot is only in hopes of dislodging new insights from the bottom that might raise your protections to the top.

The Q2 2023 report includes:

08

Network malware and exploit trends

Our Firebox network security products detect and block thousands of network and malware attacks every day. This section highlights the trending malware and network attacks opted-in Fireboxes blocked during the quarter. We share the top threats by volume, the most widespread, and by region, and we highlight interesting malware samples we found along the way. We also illustrate how malware detected in encrypted traffic trends differently than malware found in unencrypted traffic. Highlights from Q2 include a general increase in malware overall, but a stark decrease in zero day malware. That said, encrypted traffic still contains the most malware, zero day or otherwise.

15

Top Malicious Domains

Using the DNSWatch service, we share trends about the malicious web links your users click. We block your users from reaching these domains, which is good, but it's still good to know about what malicious sites attackers have made and which ones entice your users. We share the top phishing, malware, and compromised sites we blocked, and detail what some of those sites do. For instance, this quarter we noticed a legitimate baseball WordPress blog and URL shortening domain being reused for malicious purpose.

25

Endpoint malware trends

The types of malware you see at the endpoint tends to differ from what the network sees. In our endpoint section, we look at malware trends from an endpoint perspective, using data from WatchGuard EPDR and AD360. We share the most popular vectors that malware arrives from and information about the growth or decline of various malware types and families. For instance, we continue to see a decline in ransomware, and an increase in widespread malware affecting many computers. As far as delivery vector, scripts delivering malware is still most common, but down more than normal, while Windows file-based malware has increased. We also share insights about the groups spreading ransomware, as well as let you know what product features catch the most malware.

39

Timely defenses that match the evolving trends

The only reason we stir the pot is to discover new meaningful threat landscape finding, with the hope that you can use this knowledge to deploy the proper defenses. Throughout this report and in our conclusion, we share many timely security tips that will keep you safe, with and without our products.

EXECUTIVE SUMMARY

During Q2 2023, we have continued to rollout our new methodologies to reduce statistical outliers and improve our data, but this means it's still harder for us to directly compare some of our new results to historical values in past reports. Nonetheless, we continue to see some regular trends and highlights. In Q2, network malware detections were up overall, even though more sophisticated and evasive "zero day malware" was down significantly in general. That said, when arriving over encrypted connections zero day malware continued to remain high. We also noted more Linux-based malware in the Top 10 than we've seen in the past.

Network attacks (IPS detections) dropped extremely during Q2 2023, but this is mostly due to our new methodology for dropping outliers so that we can concentrate on the most common trends. While it might seem like a big change, we suspect the higher numbers in our past reports may have been similar outliers. Overall, the Top 10 network attacks still consist of older software vulnerabilities since they are easy for attackers to find exploits for. CISA recently confirm this theme of cybercriminals targeting older vulnerabilities.

Our endpoint malware products show a slightly different perspective, with overall malware detections down a bit. However, widespread malware affecting more than one machine is up. We continue to see ransomware detection declining, down 21.6 percent quarter-over-quarter (QoQ). Even though ransomware detections are down by volume, ransomware groups are still breaching and extorting many companies, and we share some examples in our report. As far as delivery, malicious scripts and living-off-the-land (LotL) techniques remain high, but have dropped some, and were replaced by more Windows-specific malware files.

The report contains a lot more detail, including information about the top malware and phishing domains we blocked users from, but we will save that for later in the report.

Below are the top executive highlights from Q2 2023:

- **This quarter, we moved to "per Firebox" malware volume reports.** Below are the malware results for our various malware detection services:
 - **Average total malware detections per Firebox: 1,177**
 - **Average malware detections by GAV per Firebox: 516** (43.8% of total malware)
 - **Average malware detections by IntelligentAV (IAV) per Firebox: 503** (42.7% of total malware)
 - **Average malware detections by APT per Firebox: 158** (13.4% of total malware).
- We extrapolate that if all the Fireboxes reporting to us had all malware detection services enabled, we would have had **88,450,373 malware detections during Q2 2023**. Note, that number only represents the Fireboxes that have opted into sharing data with us. It does show if you do not have, or have not properly configured our Total Security services, you may be missing a lot of malware.
- **Endpoint ransomware detections declined ~22%**, continuing from a 73% decrease last quarter. Despite that, heavy ransomware extortion activity remains, so keep your ransomware defense strategies current. This translates to **981 attacks blocked per 100k endpoints**, which is an 8.2% decline from Q1.
- **95.6% of malware hides behind encryption!** With a single point decrease over last quarter, but still high in general. We've mentioned it before, but most malware hides behind the SSL/TLS encryption used by secured websites. If you don't inspect this traffic, you are missing most malware your network security controls. While your endpoint malware protection acts as a safety net, we highly recommend scanning encrypted traffic.
- **Zero day malware dropped to only 11% of all malware**, which is an all-time low. We aren't sure what caused this huge decrease. However, it **remains high at 66% of malware seen over encrypted connections**. We suspect attackers primarily deliver malware over encrypted connections now.
- **Four Linux-based malware variants made our Top 10**, showing that attackers increasingly targeted Linux last quarter.
- **Ransomware detections are down on endpoints with only 465 detections per 100K endpoints**, which is a 21.6% QoQ reduction and a 72.4% year-over-year (YoY) drop.
- **Office documents and adware make the most widespread malware.** Our widespread malware list features the malware that touches the most victims, even if it's not technically the highest pure volume. We continue to see document-based threats, targeting Office products, in this list, but also had a few instances of adware on the top.
- **Network attacks dropped almost 80 percent quarter over quarter (QoQ).** However, this is more likely due to the change in methodology we now use to remove outlier data.
- **Cybercriminals continue to target older software vulnerabilities.**
- **The top 1% of Fireboxes saw 41.5% of total detections**, showing that a small minority of devices still receive the most attacks.
- **Regionally, network attacks are fairly evenly distributed across the world**, with EMEA having a tad more, and APAC receiving the least. Before our change to remove outliers there were bigger gaps in this regional distribution.
- **On endpoints, 981 attacks were blocked per 100k machines.**



- **Link-shortening domains and hijacked WordPress blogs are leveraged for malicious purposes.** When researching the most common malicious domains we blocked this quarter, we saw a new baseball-related WordPress blog and a domain-shortening service targeted.

- **Script malware delivery is down,** while other malware vectors, including Windows files, are up. While scripts like PowerShell still remain the most common malware delivery vector, they are down significantly QoQ. Meanwhile, malware leveraging Windows files is up.

That is only a fraction of data the report contains. Most importantly, each section goes into specific examples about some of the malware, threats, and attacks that you may find interesting and can learn from. Keep reading for more details about these trends and tips to protect yourself.





FIREBOX FEED STATS

WHAT IS THE FIREBOX FEED?

The Firebox Feed is our source of anonymized primary data from Firebox customers that have opted in to sharing threat detections with WatchGuard. This data 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

1,177

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

Basic Gateway AntiVirus
(GAV) service

516

Basic antivirus increased another **41%**

APT Blocker (APT)

158

Advanced evasive malware detections decreased **52%** from the previous quarter

IntelligentAV (IAV)

503

A whopping **113%** increase in IntelligentAV

GAV with TLS

763

A huge increase of **199%**

APT Blocker with TLS

844

Encrypted evasive malware dropped **15%**

TLS malware %

95.6%

Over **95%** of malware detections come from an encrypted connection

MALWARE TRENDS

Our Firebox Feed data provides us with a sample of malware traffic around the world. This real data, gathered directly from active Fireboxes, provides a useful view of past malware trends. No one can predict what new malware will appear or do, but by studying historical malware trends you can identify what threat actors have done before and how you can adjust to block it. You can leverage those lessons to avoid future mistakes. Every quarter the WatchGuard Threat Labs researches and analyzes this data, and we present our conclusions on takeaways here.

During Q2, several new Linux-based malware families made it into our Top 10 Malware table. They all target victim servers with either a botnet, cryptocurrency mining, or both. We have seen Linux-based malware droppers in the past, but this is one of the first times we have seen ones drop botnet trojans. We also haven't seen as many malware samples targeting Linux at the same time as they did this quarter. In summary, four of the top 10 malware detections in Q2 target Linux.

Before getting into those details, let's review the Q2 highlights.

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

Our Top 10 Malware list shows the malware families detected most often from the GAV and IAV services. Each Firebox detected over 1,100 malware hits on average, but these top 10 samples make up a significant portion of the detections.

We made a change this quarter. Most of the time we see very little difference in the malware samples inside malware families, but in Q2 we found a few malware families that contain malware with completely different infection paths. For this reason, we have also added a variant suffix to the end of the threat names. Further complicating this issue, one of the malware samples in the Top 10 table downloads another malware sample in the table. Linux.Generic.295484 will download Linux.Generic.13476 or what we call Linux.Lucifer. You will find our analysis of these malware samples at the end of this section.

Linux-based malware exploded in Q2 with four of the Top 10 malware families targeting Linux servers. A resurgence of XORDDoS, that we cover later, almost had the most detections of any malware family. We also found Linux.Zojfor.C.72E46613, which often uses exploits to gain access to Linux servers. For example, it leverages the Log4J exploit called [Log4Shell \(CVE-2021-44228\)](#), which made big headlines late 2021. This Linux.Zojgor variant then adds the server to its botnet and a cryptocurrency mining effort. We also saw the Zusy banking trojan with the most detections in Q2.

Threat Name	Malware Category	Count	Last Seen
Zusy.255797	Win Code Injection	654,302	Q1 2023
Linux.XORDDoS.AT	Dropper	620,905	new
Generic.3112968	Adware	2685,01	new
Ursu.808394	Dropper	116,625	Q3 2022
GenericKD.66409812	Win Code Injection	113,931	Q1 2023
Linux Zojfor.C.72E46613	Dropper	91,437	new
Logan.581	Password Stealer	87,555	Q3 2020
Linux.Generic.295484 (Linux.Lucifer)	Dropper	87,417	Q2 2022
Generic.3106131	Adware	83,674	new
Linux.Generic.13476 (Linux.Lucifer)	Coinminer	82,305	Q2 2022

Figure 1. Top 10 Basic Malware Table

Top 5 Encrypted Malware Detections

We covered the most common malware, but what about the most sneaky? Many administrators don't configure their Fireboxes to scan encrypted traffic, even though it's a free feature on all Fireboxes. We've found that 96% of malware arrives over an encrypted connection. That's why we believe our Top 5 Encrypted Malware table likely better represents the top malware in general, perhaps more so than our Top 10 list. Let's cover some of the encrypted malware Fireboxes detected in Q2.

First, in the Top 5 Encrypted Malware table, GenericKD.66409812 references an email with an attachment. That attachment attempts to download a [remote access trojan \(RAT\)](#). The malware sample we found contained two files; a PowerShell script that opens the other file, and a PNG file with an archive file appended to it. Even though an image, such as a PNG file, can't typically run by itself, they can be used to add secret data that contains portions of malware that some other process might extract (in this case using PowerShell). So, in general, you and your security controls should still watch for suspicious images like PNGs. We also saw a phishing page, the Logan password stealer, and another two code injection malware packages that target Windows.

Threat Name	Malware Category	Hits
GenericKD.66409812	Win Code Injection	113,927
Fake.Login.G.1EDAC8D2	Phishing	11,697
Heur.BZC.PZQ. Pantera.14.3C42A24C	Win Code Injection	3,423
Logan.749	Password Stealer	3,328
Heur.BZC.PZQ Pantera.14.841A73B7	Win Code Injection	2,611

Figure 2. Top 5 TLS Malware Table

Top 5 Widespread Malware Detections

The top most-detected malware list provides a good overview of the malware landscape, but each individual environment sees different threats. For example, you might expect larger, high-value targets to see more malware than small networks, or even a wider variety. That's why we also like to show a list of widespread malware families, seen by the most Fireboxes. Since the malware variants in this list hit many different networks, you should probably prioritize strategies to stop these threats, regardless of your organization size.

The most widespread malware, Adware.JS.Agent.FM, seems to target all regions equally. Adware usually just causes annoyance but can also download malware, as we saw with the Zusy adware in last quarter's report. Fireboxes also detect a lot of the Microsoft tech support scam file, Trojan.Cryxos.3903. This malware displays a page that asks the victim to call a phone number while locking the mouse and displaying fake virus popup alerts. It almost exclusively targets the US and Canada. The last three in the Top 5 Widespread Malware infected the victim through Microsoft Office exploits, delivered via malicious documents.

Top 5 Most-Widespread Malware	Top 3 Countries by %			EMEA %	APAC %	AMER %
Adware.JS.Agent.FM	Thailand - 35.8%	Malaysia - 33.07%	Indonesia - 30.85%	12.29%	10.73%	11.41%
Trojan.Cryxos.3903	United States of America - 40.92%	Canada - 17.55%	United Kingdom - 0.07%	0.01%	0.02%	31.78%
Exploit.MathType-Obfs.Gen	Greece - 27.34%	Germany - 24.24%	Hong Kong - 23.03%	16.12%	5.34%	3.70%
Exploit.RTF-ObfsObj-Dat.Gen	Greece - 28.04%	Indonesia - 22.34%	Germany - 21.19%	14.70%	5.82%	3.45%
Exploit.CVE-2017-11882.Gen	Hong Kong - 17.76%	Indonesia - 15.96%	Greece - 13.08%	9.24%	4.09%	2.26%

Figure 3. Most-Widespread Malware Table

Geographic Threats by Region

Malware targets each region differently. We saw that Trojan.Cryxos.3903 targets the US and Canada while many of the Office exploits target Europe, the Middle East and Africa (EMEA). Overall EMEA saw more hits than both other regions combined. This doesn't mean EMEA has more malware though, because per Firebox, Asia Pacific (APAC) has the most. To summarize, during Q2 Fireboxes reported 41% of detections from APAC, 38% from EMEA, and 21% from the Americas (AMER).

This regional distribution represents an increase in focus in the APAC region from previous quarters. We don't know the cause for this increase, but we could speculate it comes from heightening political tensions in the region. However, it could also amount to the normal ebbs and flows of malware in each region.




Figure 4. Geographic Threats By Region

Catching Evasive Malware

Signature-based malware detection provides split-second results on whether a file is malicious, but not all new malware has a known signature immediately. Zero day and/or evasive malware can bypass signature protection and get past legacy protections. Intelligent AV and APT Blocker malware detection services are designed to catch this more evasive malware using more sophisticated techniques beyond signatures. However, we saw a drop in detections from APT Blocker this quarter. That said, zero day malware detections over TLS connections (encrypted) stayed roughly the same.

Like past reports, our devices show more malware detection in encrypted connections, suggesting that threat actors are carrying out attacks of secure web links. Fireboxes not set up to scan encrypted traffic will likely miss two-thirds of zero day malware. Our TLS/SSL decryption feature is free with any Firebox license, so we recommend you use it.




Figure 5. Zero Day Malware

Many times, Windows-based malware infections start as a simple PowerShell command that downloads a malicious file. If the download fails, some of these malicious PowerShell scripts attempt to download another file from a different location, perhaps over an encrypted connection. This is why defense in depth is so important. You might block the first attempt at downloading malware but if you don't scan encrypted connections for zero day malware, you leave that path open.

Individual Malware Sample Analysis

Linux.XORDDoS.AT

XORDOS targets Linux-based servers running an SSH server with a brute force attack to attempt to gain access. When it succeeds and installs, it communicates with a command-and-control server (C2) to get instructions that cause the new victim server to target other Linux servers with SSH brute force attacks. Besides targeting Linux servers, XORDDOS's unique botnet uses multiple evasive and persistence techniques to stay on the infected system. The name XORDDOS comes from the malware's use of XOR for basic encryption for C2 communication and malware installation.

Finally, it attempts to find any SSH keys on the system and connects to any hostnames it finds on your system with those keys. In the table below you can see the variables named with the respective fields. A recent surge in XORDDoS.AT reveals renewed effects from



Figure 6. XORDDOS

its operators to increase the botnet size. See this [blog post](#) for more details on the botnet and the DDoS attack.

Trojan.Linux.Generic.295484- Application.Linux.Generic.13476 (Linux.Lucifer)

The nondescript name Linux.Generic.295484 seems to have connections to the [DDoS Cryptominer malware called Lucifer](#), or it could be a copycat but targeting Linux operating systems (OS) instead of Windows. We will call it Linux.Lucifer in our report.

Linux.Lucifer contains a shell script that shuts down and removes many applications running on Linux servers. It doesn't just target anti-malware or security process for shutdown, but many resource-intensive services from SQL to Python. We presume it does this to regain said resource for its own purposes. Once completed, Linux.Lucifer downloads the Monero Coinminer for Linux and applies a configuration that mines for the attacker's Monero wallet. We will get back to the configuration file in just a moment. The Monero Coinminer – though sometimes used legitimately – matches the signature for Application.Linux.Generic.13476 is also present in our Top 10 list for this quarter.

We identify the Monero Coinminer as malware even though you may have legitimate reasons to download this program outside of a company's network environment. The fact that Linux.Lucifer downloads this file confirms its use for non-legitimate purposes.

After Linux.Lucifer downloads the Coinminer, it contacts [http://18.\[130\].\[193\].\[222\]/wp-content/config.json](http://18.[130].[193].[222]/wp-content/config.json) for its configuration and starts mining for the threat actor's wallet, which you can see as the user in the config excerpt below:

```
“url”: “pool.supportxmr.com:80”,
“user”: “49VQVgmN9vYccj2tEgD7qgJPbLiGQcQ4uJx-
TRkTJUCZXRruR7HFD7keebLdYj6Bf5xZKhFKFANFxZh-
j3BCmRT9pe4NG325b.lucifer”,
```

Linux.Lucifer uses the mining pool Supportxmr and you can see the alias “Lucifer,” in the user parameter. This connects this malware to the lucifer DDoS malware on Windows systems. If one of your Linux servers is infected, it might stop doing its intended purpose and may become sluggish due to the mining.

IAV: Trojan.GenericKD.66221604

We identified malware caught by IntelligentAV (IAV) as the same malware later detected as Trojan.GenericKD.66221604. This malware attempts to download a password stealer called Lokibot. We have referenced Lokibot in the past, in our Q3 2020 ISR report. You can also learn more about it in this [CISA security advisory](#).



Figure 7. Fake DHL Icon

The variant we analyzed had the following file icon and its name indicates the file likely arrived in an DHL shipping-themed email that was meant to trick the recipient into opening it. Fortunately, IAV detected this malware before it could make it into a recipient's inbox.



Network Malware Summary

During Q2, threat actors seem to have increasingly targeted Linux servers, suggesting malware doesn't only go after Windows machines. No operating system (OS) is immune to network cyberattacks. We recommend you harden your Linux servers by not exposing SSH services if you don't need them publicly, or limiting such exposure to access control lists (ACLs) or by VPN connection. We also recommend you make sure any SSH users have strong credentials and use multi-factor authentication (MFA). Better yet, you can also set up SSH access using certificates instead of passwords. Using certificates would prevent most SSH brute force attacks.

Malware like Zojfor spreads through exploits. Because of that, we recommend keeping your OS and server software up to date, whether on Windows or Linux servers. This prevents all known and fixed exploits from affecting the server, leaving only zero day exploits (very rare) as a threat.

Finally, APT Blocker and IAV catch new and evasive malware, even when it targets Linux. Use the Fireboxes Total Security package and be sure to enable APT Blocker and IAV. These combined advanced malware detection services give you the best perimeter malware protection. That said, don't forget host-based protection solutions like Watchguard EPDR add a final layer of malware defense too.



NETWORK ATTACK TRENDS

"In 2022, malicious cyber actors exploited older software vulnerabilities more frequently than recently disclosed vulnerabilities and targeted unpatched, internet-facing systems. Proof of concept (PoC) code was publicly available for many of the software vulnerabilities or vulnerability chains, likely facilitating exploitation by a broader range of malicious cyber actors."

- 2022 Top Routinely Exploited Vulnerabilities - CISA

CISA's quote is an apt representation of the activity the WatchGuard Firebox Intrusion Prevention Service (IPS) encountered during Q2, and to a large extent, IPS activity in all past ISR reports as well. Malicious actors continue to show their tendencies to exploit old vulnerabilities as seen by CISA, WatchGuard, and the security industry at large. Threat actors exploit older software flaws because they can find exploits readily and even if most administrators patch, attackers know stragglers still exist. As CISA points out, low-cost, high-impact vulnerabilities with a long lifespan for exploitation are often the path taken by attackers. That is why IPS is an integral part of the WatchGuard Firebox service collection, as its signature-based approach protects customers from old (and new) heavily exploited vulnerabilities. It is why we see the 2021 Microsoft Exchange Server ProxyLogon vulnerability, even two years after its discovery. By the way, several CVEs (Common Vulnerability Enumeration record) related to ProxyLogon are among CISA's additional list of top exploited vulnerabilities.

While 2021 vulnerabilities are considered old, several of our new detections in Q2 block vulnerabilities from 2016 and 2017. That isn't necessarily old, based on some signatures we see for vulnerabilities dating back to 2010 and further. However, while the CISA report is focused on widespread vulnerabilities, our new Q2 detections affect old open-source software, not necessarily with a large following of active users. At least not comparable to Microsoft-related or Apache Struts vulnerabilities. Those software products, one learning management software and another for a home streaming media server software, are discussed in the sections below.

Our IPS detection volume during the second quarter was 93 detections per Firebox, which is down quarter-over-quarter (QoQ). Without context, that may seem like a large number, but the IPS activity graph (Figure 8) illustrates the huge decrease; a near 80% drop from last quarter. This is due to how we revised our methods of normalizing our data by expanding what we consider outliers, and therefore worth excluding from the report. In short, a single Firebox seeing an exponentially more detections of one threat than many other Fireboxes combine could be a false positive. To adjust for this, we drop those outliers to get more normalized attack

trends.

The decrease in detection data included changes to a lot of our historical statistical trends. The total volume per individual signature in the Top 10 is 2.10% or less, while last quarter the 10th top signatures represented 2.6% of total IPS detections. Another noticeably changed number is the total detections per Firebox by region. Each region had an average that was under 100 detections per Firebox, while just last quarter the AMER region was 804 and EMEA was 345 detections per Firebox. We hope our new methods to remove potentially anomalous data result in more accurate overall trends. However, since our numbers have changed so much,

Quarterly Trends of All IPS Hits

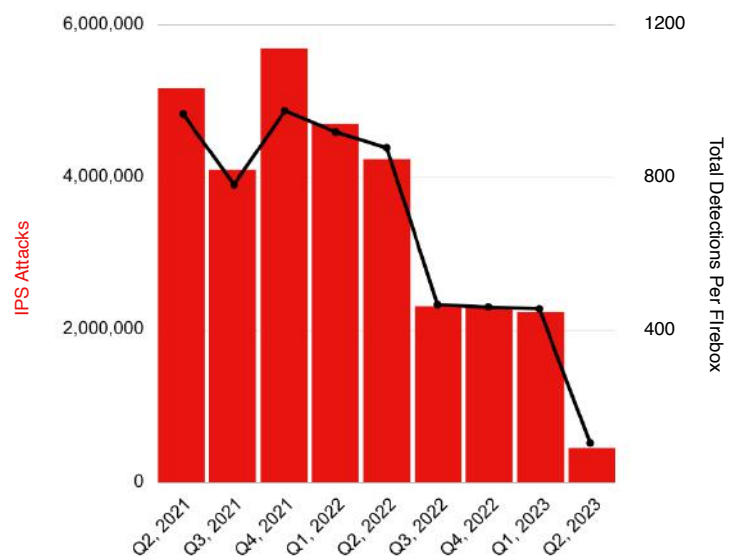


Figure 8. Quarterly Trends of IPS Hits

Unique IPS Detections




Figure 9. Unique IPS Detections

we do plan to more deeply analyze the data we are excluding to make sure we aren't throwing out the wheat for the chaff.

Additional Data:

- There were 416 unique signatures this quarter; in line with other quarters as a 3.48% increase from last quarter, and 6.97% decrease since Q2 2022. On average, we had a 1.67% increase quarter-over-quarter since Q2 2020.
- The top 1% of Fireboxes by total volume handled represents 41.5% of total detections. The top 10% represented 78.7% of traffic. Fireboxes are built to handle extreme such as what the 1% regularly handle. The concentration continues to decrease since 2021, where the top 1% used to represent 75% of total volume, and the top 10% when it was 92% of total volume.
- On average since Q2 2020, total volume has increased 0.45% quarter-over-quarter. Some quarters have had wild changes, such as an 89.96% increase for Q3 2020, to a 45.51% decrease in Q3 2022.

Top 10 Network Attacks Review

The top 10 network attacks consist of signatures with the most detections. These Top 10 signatures each represent 0.70%-2.10% of total detections. While the individual percent may seem small, that is a significant ratio when you realize that there are 416 unique signatures total detected this quarter, the remaining 406 likely representing a much smaller fractional percent individually. As we suspect there are anomalies (false positives) among our customers data, we exclude signature data if it goes beyond our statistical deviation standards. This means we have a large segment of customers who encounter these top signatures, even if it is not as common as the signatures found in our most-widespread attacks.

ProxyLogon, or Signature 1138800 has been present in a top spot of our Top 10 list since Q3 2022. Initially in 8th place, then in 4th the last two quarters, and finally reaching 1st this quarter. It isn't a surprise to see the 2021 Microsoft Exchange Server critical vulnerability, known popularly as ProxyLogon, in the top spot as any authentication bypass attack against a Microsoft Exchange Server is a juicy target. Six other signatures were present in the Top 10 last quarter, many of which have been present on this list for several years, still remain despite our new methods to drop outliers. The other three signature detections are new to Q2. Two of those were present among our Top 50 signatures last quarter, but only one, signatures 1054556, in 10th place this quarter, has no previous history among the top signatures. We will discuss that one and the other two signatures next.

Signature [1132793](#)

Signature 1132793 detects an SQL injection vulnerability in ATutor software. An additional authentication vulnerability was present as well. ATutor is an open-source learning management system (LMS); software for developing and managing online education courses. The vulnerability had some roadblocks, as it required a user to be logged into the student coursework account, which was only possible if remote registrations were enabled. Once logged-in through a student account, the attacker could bypass

authentication to reach the administrator account and inject malicious code.

This vulnerability affects version 2.2.1, from 2016. The latest version, 2.2.4, was released in 2018, the same year the owner of ATutor, Greg Gay, stepped down as its leader (and active maintainer) after 20 years of starting the project. While the website is still active and the software is available, the project is maintained at a bare minimum. Developers can continue to contribute if they so choose by submitting a pull request to fix identified issues. Only a handful of pull requests have been approved since Greg stepped down in July 2018. The last two pull requests were approved in 2019 and included fixes for a vulnerability and bug. In 2021, two separate users submitted issues, one involving an arbitrary password reset affecting version 2.2.4, and the other for an account takeover vulnerability (without details included). In both cases, Greg responded and recommended submitting a pull request to address the issue; otherwise it would not be addressed. Neither of those users decided to contribute, which is fine, as it isn't their responsibility, nor Greg's to contribute as neither are paid for their time. That, unfortunately, is the reality for products no longer maintained, for both open-source and paid software. Though it is especially common among open source projects, as unpaid maintainers simply have little incentive to commit their time once they lose passion for a project. Paid software, on the other hand, usually continues some maintenance as long as its user keep paying the provider. This is a clear example where users of a no-longer maintained product (and likely small active community) should try to shift to an alternative solution for security concerns. If that isn't an immediate option, then using the IPS solution and other security tools can at least prevent some attacks.


Signature [1132891](#)

This signature detects several integer overflow [PHP](#) vulnerabilities. Published in 2016 by PHP, these vulnerabilities are due to flaws in how PHP versions 7.x (before 7.0.6) improperly parsed ZIP files using the `getFromIndex()` and `getFromName()` methods of the `ZipArchive` class. Hans Jerry Illikainen discovered the vulnerability, and it can be found in the [Exploit Database](#).

Signature [1054556](#)

This newly discussed signature is different from the two previously discussed signatures, as it is appearing in our top 50 signatures (though we only routinely highlight the Top 10) for the first time. Signature 1132793 was in 25th place last quarter and as far back as 46th in Q1, 2021. Signature 1132891 has only been present since last quarter when it was in 28th place.

This signature goes back to 2010 when a buffer overflow vulnerability was found in HP OpenView Network Node Manager (OV NNM). A program within the product failed to handle HTTP requests if the value sent to the mapping graphic application was beyond the parameter limits (classic overflow). In addition, exploiting this required overwriting the structure exception handler (SEH) in the memory stack, used for addressing errors. In other words, this is a classic stack buffer overflow vulnerability, one



of the most basic of many types of memory corruption vulnerabilities.

As the OV NNM name suggests, it is a network monitoring software for tracking organization servers and other assets. It was the first of many products within the OpenView suite. HP OV NNM is an old product, rebranded several times, and then in 2016 HP transferred their Software

Signature	Type	Name	Affected OS	Percentage
1138800	Web Attacks	WEB Microsoft Exchange Server Remote Code Execution Vulnerability -6 (CVE-2021-26855)	Windows	2.10%
1059877	Access Control	WEB Directory Traversal -8	Windows, Linux, FreeBSD, Solaris, Other Unix	1.60%
1132092	Buffer Overflow	FILE Invalid XML Version -2	Windows	1.40%
1055396	Web Attacks	WEB Cross-site Scripting -9	Windows, Linux, FreeBSD, Solaris, Other Unix, Network Device	1.40%
1059958	Web Attacks	WEB Directory Traversal -27	Windows	1.30%
1132793	Web Attacks	WEB SQL injection select from attempt -5.a	Windows, Linux, FreeBSD, Solaris, Other Unix, Mac OS	1.30%
1132891	Buffer Overflow	WEB PHP ZipArchive getFromIndex and getFromName Integer Overflow (CVE-2016-3078)	Windows, Linux, FreeBSD, Other Unix	1.20%
1054837	Web Attacks	WEB Remote File Inclusion /etc/passwd	Windows, Linux, FreeBSD, Solaris, Other Unix	0.80%
1056773	Buffer Overflow	WEB Web Server Connection Header Buffer Overflow	Windows	0.80%
1054556	Buffer Overflow	EXPLOIT HP OpenView NNM ovwebsn-mpsrv.exe Invalid Option buffer overflow (CVE-2010-1960)	Windows	0.70%

Figure 10. Top 10 Network Attacks by Volume

Top 10 History




Figure 11. History of Prominent Signatures in the Top 10 Since Q2 2020

Business Segment assets (including OV NMM) to Micro Focus. The “Network Node Manager i” at Micro Focus is the latest iteration of this software, spanning over three decades.

The above Top 10 History chart shows how common it is for the same signatures to remain near the top of the list. Likewise, as we have discussed the Top 50 signatures, many of those present in the Top 50 remain there, with only an average of five new signatures appearing each quarter. This is simply a way to show how many vulnerabilities remain popular targets. Signatures 1138800 (in yellow) managed to reach the top spot because of how significant it would be to compromise a Microsoft Exchange Server.

New Signatures in the Top 50

Each quarter, we grab a list of the top 50 IPS signatures by volume. Of that list, we present the Top 10 signatures. In addition, we use the full Top 50 list to gain insight when signature detections shift widely quarter-over-quarter. As it is very common to see a cycle of 10–15 signatures make its way in and out of the Top 10 signatures, we thought it would be interesting to learn about what signatures are entirely new to the top detections list. That is, a signature is considered a ‘new top detection’ when it has never previously appeared in the Top 50 signatures by volume. These signatures are often present in the signature database for years, but it is possible for the signature to be relatively new in terms of published date as well. This quarter we had six new signatures detection. A normal count, as the average has been five signatures per quarter since 2020. One of those we already discussed is signature 1054556, in 10th place. The others are described below with a few additional details.

Signature [1231780](#)

In 17th place, this vulnerability affected HAProxy before 2.7.3. It is a popular open-source software load-balancer and reverse proxy. Red Hat discovered this vulnerability in February 2022 and worked with the HAProxy teams to remedy the issue. The vulnerability in question impacted the Native HTTP Representation (HTX) which was introduced in HAProxy 1.9 and included by default in all subsequent versions of their latest 2.0 software. HTX is HAProxy’s own solution for maintaining consistent and high-performance operations when handling all current and future HTTP protocols. The vulnerability was from a failure to properly handle HTTP responses containing the “Set-Cookie2” header, resulting in infinite loops.

Signature [1230310](#)

This is the only denial of service (DoS) or distributed denial of service (DDoS) type attack in the Top 50. It is more widely known to security industry professionals as the Slowloris DoS attack tool. While the signature was updated in 2021, this is an old vulnerability from 2009 when the tool was released. Since then, there are now a myriad of mitigations in place to prevent attacks from Slowloris and other DoS/DDoS-based attacks. Firewalls and reverse proxies are examples of tools that can mitigate these attacks. DDoS attack methods continue to evolve and at scales once unthinkable in

2009. We hear repeatedly in the news about record-breaking DDoS attacks, but it is important to be cognizant of attacks targeting smaller organizations as well.

Signature [1133728](#)

This is a directory traversal vulnerability in the Trend Micro Threat Discovery Appliance 2.6.1062r1. Researchers discovered this in 2016 and [published a proof of concept](#). The exploit is accomplished through authentication bypass by deleting a config to reset back to the default admin password and waiting for a server reboot. The researchers mention that the deletion of the config file would likely trigger a needed restart of the servers anyway as the deletion of the password file would prevent other users from logging in.

Signature [1133500](#)


This signature encompasses various Apache Struts 2 vulnerabilities with CVEs from 2016 and one from 2017. Apache Struts is an open-source framework for building Java applications. A remote code execution (RCE) can be performed based on numerous issues. Those from 2016 are: “XSLResult” function, “method:” prefix, and the REST Plugin. The function “XSLResult” converts Extensible Stylesheet Language Transformations (XSLT) to XML, which failed to properly validate the type and content from files. The “method:” prefix when Dynamic Method Invocation (DMI) was enabled. The REST Plugin is vulnerable for several reasons, one when using the “!” operator with the plugin. The solution is to either disable DMI, or upgrade Apache Struts versions.

The critical 2017 vulnerability, [CVE-2017-5638](#), garnered the greatest attention. The Jakarta Multipart parser used in Apache Struts was being exploited in the wild.

Attackers triggered the RCE by sending maliciously crafted Object Graph Navigation Language (OGNL) expressions in “Content-Type” values. Many months later, another two more RCE vectors were found by sending malicious “Content-Disposition” values, or by having improper “Content-Length” headers. Authentication wasn’t necessary to carry out this attack. The ease of exploit is especially problematic, as there were numerous victim web applications being exploited days after the patch was released. This isn’t surprising, as pushing updates to Struts-based applications can be a slog due to the need to rebuild scripts and update dependencies. There are numerous openings for the applications maintainers to run into issues, and that’s just for one application. The number of Struts-based applications in an organization could be extensive. Therefore, any network defensive tools (like IPS) can be crucial in scenarios like these. Even six years later, there may be Apache Struts applications still vulnerable to this exploit.

Signature [1135666](#)

This signature is for a vulnerability in PHPUnit (before 4.8.28 and 5.x before 5.6.3), a testing framework for PHP. The 2017 buffer overflow



Signature	Type	Name	Affected OS	Rank
1054556	Buffer Overflow	EXPLOIT HP OpenView NNM ovwebsn-mpsrv.exe Invalid Option buffer overflow (CVE-2010-1960)	Windows	10
1231780	Web Attacks	WEB HAProxy h1_headers_to_hdr_list Empty Header Name Access Control Bypass (CVE-2022-0711)	Network Device	17
1230310	Dos/DDoS	WEB Slowloris Tool HTTP Denial Of Service -1 state2-F/flow	Windows, Linux, FreeBSD, Other Unix	34
1133728	Web Attacks	WEB Directory Traversal in Cookies	Windows, Linux, Other Unix	43
1133500	Web Attacks	WEB Apache Struts Dynamic Method Invocation Remote Code Execution -1.u	Windows, Linux, FreeBSD, Other Unix, Mac OS	45
1135666	Buffer Overflow	WEB PHPUnit CVE-2017-9841 Arbitrary Code Execution Vulnerability	Linux, FreeBSD, Other Unix	49

Figure 12. New Signatures This Quarter

Most-Widespread Network Attacks

The most-widespread network attacks, as the name suggests, are the ones detected by the most individual or unique Fireboxes. Some top attacks may make the list just because they hit a relatively few Fireboxes in high volume. These widespread ones affect the most customer devices. The table includes a section for top 3 countries and each region. The percentages represent the proportion of customers who encountered attacks corresponding to those signatures. These signatures include several we saw last quarter as well as two new ones. Signature 1059877, a web directory traversal attack against several management softwares, is 3rd in most-widespread, and #2 in the Top 10. It has maintained a spot in the top five most-widespread signatures since Q2 2022. This signature has remained on the Top 10 list since Q4 2020, progressively rising from 10th place to 2nd this quarter. We discussed this signature in previous reports, with the unchanged conclusion that management software’s remain a prime target for attackers. We’ve already highlighted signatures 1130592 and 1110932 in previous reports, so we will maintain our focus on the two new signatures this quarter.

Signature [1133215](#)

This signature consists of two CVE’s from Microsoft for ActiveX address browser memory corruption vulnerabilities in their Internet Explorer and Edge browsers. ActiveX is a Microsoft framework used in a wide-range of Windows products. This has since been deprecated. Exploitation is only feasible if the attacker can entice a user to a domain embedded with malicious code. Commonly, that occurs via a phish leading to a newly created malicious domain, or a compromised domain. The attacker only gains the privileges of the victim’s Windows user, but if they have local system administrator rights, as many Windows users do, this could result in computer control of that Windows machine.

Signature [1134586](#)

This signature detects XML external entity (XXE) processing vulnerabilities in two types of software. One is for Subsonic, a media streaming server software. The other affects Microsoft Common Console Document (.msc) in Microsoft Windows 7 SP1, Windows Server 2008 SP2 and R2 SP1. The “.msc” is part of Microsoft Management Console (MMC) that is used by system administrators and others for managing systems. XML document parsers may deal with Document Type Definition (DTD), which is used for determining external entities such as XML and HTML. In this case, malicious actors can use a substitution string in the Uniform Resource Identifier (URI) to reach an unintended resource. The URI is a path to a resource, so a reference to an external entity with a “file:///” URI can search for contents on the users local file system and return the values as output. If the attacker sets the URI to a “http:///” path, they could then call out to an external server, possibly evading detection of exfiltrated contents. This vulnerability is a legitimate risk, but a subdued one for the Microsoft products as it was privately reported to them, and updates were pushed out. The same cannot be said of the Subsonic software.

We can presume this signature is heavily skewed to the Microsoft-related vulnerability. But Subsonic is worth touching on as it falls under the same category as the ATutor software previously discussed in the Top 10 signatures. Both are projects without maintainers. The only difference is that Subsonic went closed-source after version 6.0-beta1. The Subsonic website is still up and looks reasonably modern. A person without much knowledge of the product, but looking for a music streaming server solution, could easily find themselves on this website. The downloads are still available, and the option to pay a monthly or lifetime subscription is available. But if you go to the change log, you will find the latest update was in November 2019 for a bug fix. Beyond a few other updates in 2018, 2017 was the last year of consistent development. The forums confirm that Subsonic is essentially “abandonware”.

Even so, there is still a reasonable amount of user activity on the forum. Therefore, this is software very much in use. While there are open-source alternatives forked from Subsonic's pre closed-source code, a notable one being Airsonic-Advanced, the process for installing the software isn't necessarily straightforward. To download and manage an open-source software from GitHub is not simple. A quote from one of the forums captured it well, "I'm too dumb to get Airsonic working." I think many would concur. Therefore, many people continue opting to download software that is no longer maintained and likely full of vulnerabilities. Especially as this software allows for outbound connections from the home network to stream the music.




Figure 13. Countries listed among one or more widespread attack signatures who were most affected

Network Attacks by Region

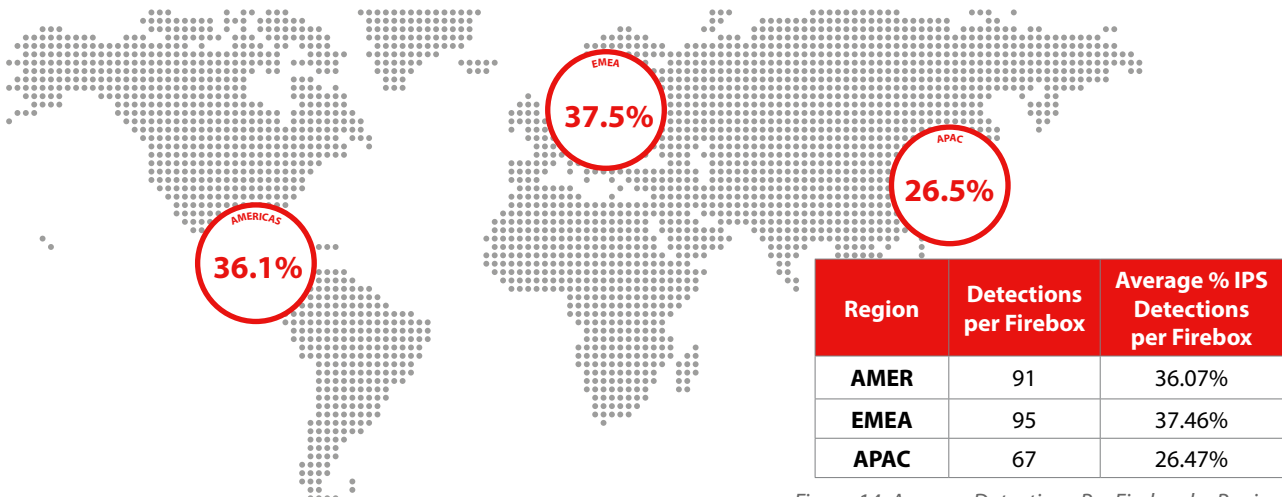


Figure 14. Average Detections Per Firebox by Region

Finally, let's look at network attacks on a regional basis. The detections are often disproportionate between the regions, so we seek to normalize this data. Figure 16 shows average detection by Firebox per Region. Several things stand out. The wild variances between regions began to drop after Q2 2022, at least at a more reasonable scale. The significant gap between regions in addition to total decrease in overall volume between them was attributed several top-heavy signatures. Many signatures in the Top 10 contained 15-33% of total signature volume. The numbers are more balanced this past year.

The decline in detections, overall and per region this quarter is obvious. Last quarter AMER had 804 detections and APAC had 286. EMEA was 95 this quarter, which is quite a drop for that region as it has had a relatively stable detection range since Q3 2021. The APAC region pattern continued to climb between Q2 2021 and Q2 2022. As we mentioned earlier, the drop in detections, especially for AMER and APAC, have been attributed to the top-heavy signatures. The actual reason for the large drop in numbers this quarter is due to how we updated the exclusion of our outlier data. We must have had several Fireboxes producing an extraordinary amount of detection above the average in particular regions. Therefore, the detections are now sub-one-hundred after being in the multi-hundred range last quarter. Next quarter we will be able to gage if this is a new normal.

Average Per Firebox Detections by Region




Figure 15. Average Detections per Firebox by Region since Q2 2021

The detections percent by region show another way to view the regional numbers. Both AMER and EMEA were nearly on par while APAC was smaller, but not significantly. Were we to show our raw data without considering the number of Fireboxes by region enrolled, EMEA would have contained a larger share of detections. By considering several other factors, the percentage by region (as well as detections per Firebox) show the average reality for Fireboxes in each region.

Detections Percentage by Region




Figure 16. Average Detection per Firebox Percentage since Q2 2021

Conclusion

There were a lot of noticeable changes this quarter in terms of raw data. By updating our data specifications to exclude outlier Firebox detections, we improved the quality of data relayed into this report. Subsequently, some of our graphics such as IPS activity, and average detections per Firebox by region, show a steep drop. As we continue publishing these reports, the data should begin to level off and begin to form new patterns that we can analyze.

A reoccurring theme for the IPS section is the continued prominence of old vulnerabilities. As CISA mentioned in their report on top exploited vulnerabilities in 2022, malicious actors often seek the easiest path towards exploitation. That means taking advantage of vulnerabilities that they know are widespread and easy to exploit, even if that means using attacks that are several years known. We know, CISA knows, and attackers know that it's a numbers game. Hence why the top signature this quarter is a 2021 vulnerability attempting to exploit Microsoft Exchange Servers. Even so, there are several old vulnerabilities targeting very specific software and with a likely small userbase. ATutor is one of them. Supersonic is another. Both types of software are no longer supported and yet they continue to have users download and use their software. Therefore, when products such as these, sometimes referred to as 'abandonware', can no longer rely on security updates, products like IPS are available to mitigate any potential attacks. That said, please migrate to software solutions that are at least pushing out security updates, even if the product enhancements have stopped.

DNS ANALYSIS

DNS-based filtering is an important tool in the defender’s toolchest for protecting against phishing and other email-based threats. Without controls like DNS filtering, if a user misses a misspelled domain name, or is tricked into visiting a compromised website, they stand the risk of becoming victim to modern authentication attacks that can even circumvent MFA. Tools like DNSWatch act as that last line of defense when the user clicks, by redirecting them to a secure black hole instead of the original malicious destination.

In this section, we cover the top malicious domains that DNSWatch protected users from visiting in the quarter, bucketed into three main categories: malware domains, phishing domains, and compromised websites.

Top Malware Domains

Domain detections in this category include the domains and websites attackers use to distribute malware or facilitate command and control communications. In general, these domains have no legitimate purpose and were deployed specifically for use in malware infections.

Malware
x-vpn[.]jug
greenwidow[.]top *
profetestruec[.]net *
pixel-install[.]me *
xrass[.]com
toknowall[.]com
t[.]jamxny[.]com *
hrtests[.]ru
newage[.]newminer-sage[.]com
newage[.]radnew-age[.]com

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.

Figure 17. Top Malware Domains

There were four new domains in the top blocked domains this quarter that we had never seen in the top ten list before. The first new entry, greenwidow[.]top, serves as a command and control domain for a JavaScript trojan. This trojan pretends to be a PDF attachment but instead executes a JavaScript downloader when the user clicks on it. The second domain, profetestuec[.]net, was originally added to our feed nearly three years ago after identifying it as being used by the WannaMine coinminer malware. The WannaMine malware initially starts with a malicious .bat script file that launches PowerShell. The PowerShell script downloads several modules including Mimikatz for stealing Windows account infor-



Figure 18. ViperSoftX PowerShell Loader

mation and EternalBlue shellcode for proliferating over SMB before registering the system to mine cryptocurrency for the attacker.

The next domain, pixel-install[.]me, is a communication domain for another JavaScript-based trojan associated with the Cryxos family, which we’ve discussed in the Firebox malware section in previous reports.

The final new domain, t[.]jamxny[.]com, joined our feed two years ago after a third-party source reported it for hosting the Lemon Duck malware. This particular Lemon Duck campaign appears closely associated with previous top malware domain entries in recent quarters.

Top Compromised Domains

We classify a domain as compromised when we believe it is a legitimate destination that a threat actor has corrupted to host malicious content. As an example, cybercriminals regularly compromise vulnerable WordPress websites and hide malware delivery or phishing campaigns. If a compromised site administrator doesn’t notice the malicious activity, the hosted content can remain active for an extended amount of time. By targeting and corrupting legitimate websites, threat actors can benefit from the existing good reputation of the site and evade many reputation-based defenses.

Compromised
d[.]zaix[.]ru
www[.]sharebutton[.]co
ssp[.]adriver[.]ru
granerx[.]com
dodgersdigest[.]com *
dinatds[.]com
joinmy[.]site *
a[.]pomf[.]cat
ozcontests[.]com *
fortnitechat[.]site *

Figure 19. Top Compromised Domains

Four domains in the top compromise domains list were brand new to this quarter. The first domain, dodgersdigest[.]com, appears to be a WordPress blog dedicated to news and trivia involving the LA Dodgers baseball team. Several contact pages on the website have been compromised to host spoofed Twitter feeds (complete with a 2014 Twitter copyright) that contain a form to “Sign up for Twitter” with an email and password, likely as an attempt to capture credentials for credential stuffing attacks.



Figure 20. dodgersdigest[.]com

The second domain, joinmy[.]site, is less of a compromised domain and more of a legitimate domain that has been abused for malicious purposes. This domain is a part of a URL Shortener service that threat actors have co-opted to hide a malicious destination from unsuspecting victims.

The next domain, ozcontests[.]com, appears to be a website for an educational contest, such as math Olympiads, in the Southeast Asia and ANZ regions. Prior to their takedown this year, the Qakbot threat actors had compromised this domain to host a command and control infrastructure for their botnet. The final domain, fortnitechat[.]site, is another URL Shortener domain that we found hosting multiple categories of malware including spyware and reconnaissance tools.

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 *
haxbyq[.]com
data[.]lover-blog-kiwi[.]com
t[.]go[.]rac[.]co[.]uk
e[.]targito[.]com
edusoantwerpen-my[.]sharepoint[.]com
shbzek[.]com *
bestsports-stream[.]com
mail[.]cuchost[.]com *

Figure 21. Top Phishing Domains

Three domains were new to the top phishing domains list this quarter. Two new domains, ulmoyc[.]com and shbzek[.]com, were found hosting phishing domains targeting users in India. We originally added these domains in February 2023 after finding them hosting content that attempted to mimic popular brands for SEO poisoning.

We added the final new domain, mail[.]cuchost[.]com, to our threat feed three years ago after finding it hosting a phishing campaign that mimicked Outlook Web Access. While the specific phish is no longer there, the domain remains active.




Figure 22. Fake Microsoft 365 web app

Conclusion

JavaScript-based malware continues to be a popular avenue for threat actors, largely because it isn't typically a file type that administrators can block outright on their networks due to its near-requirement for modern websites. That said, blocking JavaScript as an email attachment is a good option to limit some of your exposure. Regardless, stay on the lookout for JavaScript and other living-off-the-land techniques and deploy a layered defense that includes protections at the DNS level to keep your organization safe.



ENDPOINT THREAT TRENDS



This is only the second iteration of the Endpoint section since we've made sweeping changes to the layout and data we collect. As such, we can start to pick up patterns within the data. However, as the saying goes, "Once is happenstance, twice a coincidence, three times a pattern." Therefore, we can't start determining patterns yet but theorize what these patterns might look like going forward based on the data we've derived thus far.

Since this is the second iteration of our revamp, there aren't many superficial changes from the quarter prior. However, in a testament to continuous improvement, we've made minor tweaks to enhance understanding in some areas. The most obvious change is the additional data we've provided showing the differences between this quarter and last. The other noticeable change is removing the "New Ransomware" information and the corresponding table. We removed this because it was a data point many would say "cool" to and move on. In other words, we didn't find it particularly meaningful.

We want the data to be more meaningful, and because of that, we've bundled all of the new ransomware groups into the current active extortion groups list and removed all of the ransomware that doesn't currently have a double extortion operation. Doing this removes all of the meaningless ransomware that doesn't have any tangible risk to organizations. If they did, we would highlight these external to the active extortion groups. We hope the changes help readers better understand the more harmful ransomware groups instead of showcasing the less important ones.

MALWARE FREQUENCY

Like last quarter, we begin by discussing the overall malware frequency detected by our advanced endpoint security solution – Panda Adaptive Defense 360 (AD360) and WatchGuard EPDR. In Q1, we started displaying some of our data points as "per one hundred thousand (100k) machines." In that quarter, AD360 and EPDR blocked 1,068 attacks per 100k machines. This quarter, we observed an 8.2% reduced frequency of attacks, resulting in 981 attacks blocked per 100k machines.

As stated, uncovering any patterns within the data will take one more quarter. If the malware frequency decreases into Q3, we will draw more concrete conclusions. For example, we can discover if a roughly 8% swing in frequency is typical or relatively significant. We can also try to theorize patterns if there is another decrease from Q2 to Q3. On the contrary, if the frequency recovers to levels similar to Q1, we can theorize that the frequency is stable but still high. We can then extract real-world reasons why this might be the case.



Alerts by Number of Machines Affected

Next, we move on to the number of machines affected by each malware detection. This data is most beneficial in understanding the difference between targeted malware and those sent in widespread malware campaigns. For example, one of the malware on our Top 10 Malware for this quarter is Glupteba. Threat actors often disseminate Glupteba via phishing, malvertising, and trojanized software. These attacks, unfortunately, lead to multiple victims from the same malware – the same hash. Each victim of that Glupteba campaign adds to the number of machines affected per each alert. Below, we define and describe the parameters for which we log this data:

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

In Q2, we suspected a reduced number in all categories compared to Q1, considering an 8.2% decrease in overall malware frequency. However, this was not the case. We did see a modest reduction in alerts (-8.8%) that only affected one machine. However, to our surprise, we saw increased alerts for all other categories. We observed a slight increase in alerts that appeared on between two and five machines (6.1%), five and ten machines (6.5%), and 50 to 100 machines (4.7%). We recorded a significant increase in alerts that appeared on ten to 50 machines (22.6%) and more than 100 machines (21.5%). Considering the reduced malware frequency, this data tells us widespread malware campaigns increased from Q1 to Q2.

Alerts by Number of Machines Affected




Figure 23. Alerts by Number of Machines Affected


Number of Machines	Q1 Alerts	Q2 Alerts	Difference from Q1
1	131,279	119,735	-11,544
>= 2 & < 5	13,935	14,781	846
>= 5 & < 10	2,888	3,077	189
>= 10 & < 50	1,720	2,109	389
>= 50 & < 100	211	221	10
>= 100	195	237	42

Figure 24. Alerts by Number of Machines Affected (Table)

Alerts by Top 30 Countries Affected

This subsection aims to show how Panda’s AD360 solution has a worldwide presence. We do this by defining an Alert Coefficient – a ratio of active Panda AD360 licenses and total alert counts for the quarter – and extracting the top thirty countries. For example, if we had ten active AD360 licenses in the United States and logged ten total alerts in the United States for this quarter, the Alert Coefficient would be one (ten alerts/ten licenses = 1).

The biggest mover from the quarter prior is Cuba, which ranked 19th in Q1 and is now the country with the highest Alert Coefficient (1.84). Jordan (1.13), Malawi (1.00), Laos (0.58), and Pakistan (0.51) round out the top five. The country that moved down the rankings the most was Armenia (0.09), ranking 30th and tied with Venezuela (0.09) and Indonesia (0.09), but was 20th last quarter. However, this is good since a lower Alert Coefficient means fewer malware attacks per license. There are four new countries to the rankings – Cayman Islands (0.35), ranked eighth; Angola (0.29), ranked ninth; India (0.15), ranked 20th; and Venezuela (0.09), tied for 30th.



Country	Alert Coefficient
Cuba	1.84
Jordan	1.13
Malawi	1.00
Laos	0.58
Pakistan	0.51
Micronesia	0.50
Morocco	0.41
Cayman Islands	0.35
Angola	0.29
Mozambique	0.28
Sao Tome and Principe	0.25
Grenada	0.25
Kenya	0.24
Vietnam	0.24
Bosnia and Herzegovina	0.22
Bolivia	0.21
Bangladesh	0.20
Nigeria	0.19
Turkey	0.15
India	0.15
Macedonia	0.14
Guatemala	0.13
United Arab Emirates	0.13
Botswana	0.12
Paraguay	0.11
Singapore	0.11
Andorra	0.11
Indonesia	0.09
Venezuela	0.09
Armenia	0.09

Figure 25. : Alerts by Top 30 Countries Affected (Table)

TOP MALWARE AND PUPS

Previously, we've discussed the overall malware frequency and how this frequency translates to the number of machines affected by that frequency. We've also shown which countries are most affected based on a ratio of detections and AD360 licenses. The Top 10 Malware and PUPs defined in this section are the hashes with the most quarterly detections. These are the files that contribute the most to the overall malware frequency.

We aggregate these numbers and normalize them using our metric of detections per 100k machines. As you can assume, the ten hashes with the most detections are on our Top 10 lists. We go even further by attempting to attribute each hash to its associated behavior and, in some cases, can make a definitive attribution to a malware family. We begin with malware.

Top 10 Most Prevalent Malware

Four malware families reemerged in the Top 10 Malware list that also appeared in Q1 – Glupteba, MyloBot, GuLoader, and the EICAR Test File. However, the other six are new or unknown malware families. The ninth-ranked file in the list is a downloader for the Ammyy Admin remote administrative tool. Threat actors use this tool to perform remote actions like AnyDesk and TeamViewer. The most prevalent malware in the rankings was the trojanized 3CX desktop application.

The 3CX breach was a double supply chain compromise where attackers attributed to North Korea modified software called X_TRADER from Trading Technologies. A 3CX employee then installed this installer, executed it, and the attackers ultimately tampered with the 3CX installer. Thus, when users of 3CX began downloading and installing their software, they were affected. The top-ranking malware in the list is the trojanized 3CX installer. This attack shows the breadth of those involved. We observed 479 detections per 100k machines from users using AD360 – a staggering 48% of all detections within the Top 10 rankings.

Last quarter, there were multiple hashes attributed to the same malware family. Glupteba appeared four times, and Snake appeared twice. In Q2, there were no duplicates. Glupteba, MyloBot, GuLoader, and the EICAR Test File appeared once. We already discussed two new malware in this list – the trojanized 3CX installer and the Ammyy Admin downloader.

Meanwhile, we couldn't attribute the other four to a specific malware family. Two were droppers, one was an injector, and the tenth-ranked file was a downloader. We have provided a more detailed description of each malware below.

Once a user downloads and opens this attachment, the embedded GuLoader stealthily downloads additional malware from a remote command and control server (C2). Increasingly, these C2s are trusted sources such as Discord, DropBox, Telegram, and many others. The seventh file in the list is an unnamed information stealer and spyware masquerading as SysInfo.

For a better understanding, below is a short description of each malware classification:

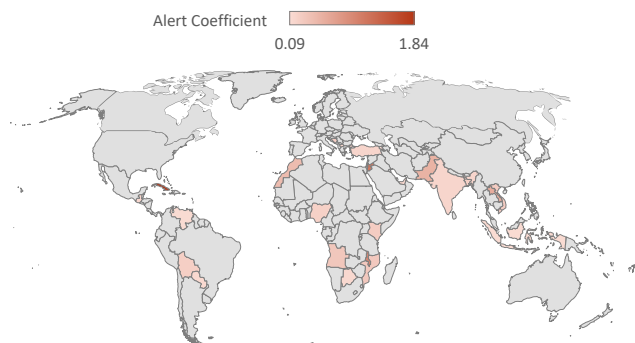


Figure 26. Alerts by Top 30 Countries Affected (Map)

MD5	Signature	Affected Machines per 100k	Classification Attestation
9833A4779B69B38E3E51F04E395674C6	Trj/RnkBend.A	479	Trojanized 3CX Desktop Application
6CC8D5F1CB1819791E4897F902FAF365	Trj/RnkBend.A	110	Glupteba
4923F1C3597619639DB2F13DB0CA44F2	Trj/Agent.OOW	95	Unknown Malware (Dropper)
2253836BB8B0B5479A1F77974B82B1F0	Trj/RnkBend.A	75	Unknown Malware (Injector)
44D88612FEA8A8F36DE82E1278ABB02F	EICAR-AV-TEST-FILE	56	EICAR Test File
3E86685246C1FDCC9EEF8B95986BA4E4	Trj/WLT.F	55	MyloBot Delivering Khalesi
91E11E5375B0D71366A26393C3C573CB	Trj/Agent.RP	35	GuLoader
A3882DD6DE6E0321FBCC171C80E8F659	Trj/CI.A	35	Unknown Malware (Dropper)
E72B313D807A536D45B68E52C1257996	Trj/CI.A	32	Ammy Admin Downloader
52CBFED702193577BCBC61E20B0B4B2C	Trj/Agent	31	Unknown Malware (Downloader)

Figure 27. Top 10 Most Prevalent Malware

Trojanized 3CX Desktop Application

In late March 2023, at the very end of Q1, threat actors breached the network of 3CX, an organization that creates communications solutions to replace traditional Printed Circuit Boards (PCBs). This breach resulted in a trojanized version of one of their applications, leading to a supply chain attack on any organization that used this software. One of the hashes associated with this attack is our number one most prevalent malware for Q2.

Glupteba

Glupteba is a multi-faceted loader, botnet, information stealer, cryptominer, 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.

Unknown Malware (Dropper)

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. In this case, a dropper is 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.

Unknown Malware (Injector)

Like before, this is a sample we cannot directly attribute to a particular family, but generically service a specific purpose. An injector is a 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.

EICAR Test File

An EICAR file, also called an EICAR string or EICAR signature, is a specific string found within a file that helps users determine if antivirus is functioning correctly. EICAR stands for the European Institute for Computer Anti-Virus Research. They developed this

standard and string with the help of the Computer Anti-Virus Research Organization (CARO). How it works is simple. If you download the EICAR test file onto your machine, your antivirus should alert you that this is an EICAR test. If it does, it means your antivirus is appropriately installed and functioning. If it does not, it means your antivirus is off, misconfigured, or uninstalled

The EICAR string:

X5O!P%@AP[4\PZX54(P^7CC)7}\$EICAR-STANDARD-ANTIVIRUS-TEST-FILE!\$H+H*

The official EICAR test file download page is here:

<https://www.eicar.org/download-anti-malware-testfile/>

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.

Ammy Admin Downloader

Threat actors sometimes use payloads that download a remote administrative access tool called Ammy Admin to perform remote commands on a victim's machine.

Unknown Malware (Downloader)

Again, a sample we can't place in a family, but generically services and exploit chain purpose. A downloader is a malware that "downloads" other malware, which seems very much like a dropper at a high-level, but only differs in how it delivers the second stage malware payload. For example, when malware uses PowerShell's Web Client to download a file from a URL.

Top 10 Most Prevalent PUPs

A PUP is an acronym for potentially unwanted programs. You may also commonly see these as PUAs or potentially unwanted applications. These terms describe 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, software installers often bundle their installation wizards with pre-selected software from third parties. The user often doesn't realize they are installing this extra software. Therefore, we assume this software is unwanted because its installation relies on deception.

Three adware-related files were on the Top 10 list with the PUP/BundleOffer classification. The second-ranked file was a MEmu Setup Software Development Kit (SDK). The installer included third-party software; thus, we classified this with the PUP/BundleOffer signature. The sixth and tenth-ranked files were installers with third-party software, too. The sixth-ranked file was an installer for Daemon Tools Lite, and the tenth-ranked file was a software called LDPlayer, an Android Emulator used for gaming.

One signature appeared four times in the Top 10 – PUP/DownloadAssistant. Usually, we reserve this signature for an application of the same name – Download Assistant. However, we also use this signature for software that facilitates the downloading of other software. Two PUP/DownloadAssistant signatures were Softonic bundle installers, one for Roblox and the other for TinyTask. These installers included third-party adware. The other two were ranked eighth and ninth and were a Vuze BitTorrent Client Installer and Praat Download Manager, respectively.

The only signature that appeared this quarter and last was HackingTool/AutoKMS. AutoKMS software are those that activate Windows products – usually illegally – or facilitate the activation of Windows products. Again, usually illegally, but not always. It only appeared once this quarter, as opposed to three times in Q1. However, the number one most prevalent PUP in Q2 was AutoKMS software that activated licenses for Microsoft Office 2013-2019.

That leaves us with the fifth and seventh-ranked files. Number five was an application for BitTorrent titled WebHelper. We classified this detection as PUP/uTorrentWeb, a classification for web-related torrent software. The seventh-ranked file was an uninstaller tool for Windows drivers titled "DriverPackSolution." We denoted this file with the PUP/DriverPack classification because it facilitated installing or uninstalling a bundle of Windows drivers. You can read more about the signature definitions below.

MD5	Signature	Affected Machines per 100k	Classification Attestation
CC470D06E9AFC9A7C0B395274B02AC88	HackingTool/AutoKMS	296	License activation tool for Microsoft Office 2013-2019
581DA0F19EF8388A0BA331CE0A617AAF	PUP/BundleOffer	174	MEmu Setup Abroad SDK
4E60FBFB9F6C7E9FE6935437253038EB	PUP/DownloadAssistant	133	Softonic Bundle Installer (Roblox)
D0DAFC349ED205185E9C30382209C1C6	PUP/DownloadAssistant	132	Softonic Bundle Installer (TinyTask)
CC70A40EEA5375C967813F0B3595B61D	PUP/uTorrentWeb	123	WebHelper tool for BitTorrent
4AE0D57D871A8D99D8340D268A23B518	PUP/BundleOffer	118	Daemon Tools Lite Installer
11FBD8034E2C62A0B5DBE718CAC49096	PUP/DriverPack	107	Windows Driver Uninstaller tool - "DriverPackSolution"
037D91C5C06601B3D6EAB400EF72157E	PUP/DownloadAssistant	93	Vuze BitTorrent Client Installer
99A9FBD5FEE72CE51585309390A46717	PUP/DownloadAssistant	90	Praat Download Manager
90276982CC921F646F74F8310EF8CD6A	PUP/BundleOffer	89	LDPlayer Android Emulator

Figure 28. 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/DownloadAssistant

Panda analysts usually use this classification for the Download Assistant application. However, files that facilitate the download of other software, such as download manager tools, are often classified as PUP/DownloadAssistant.

PUP/uTorrentWeb

These are web-related files that use the BitTorrent network.

PUP/DriverPack

Files with this classification are usually a bundle of files that are installers, modifiers, or uninstallers for operating system drivers.

Defense in Depth

The Defense in Depth subsection highlights the efficacy of AD360 and how implementing multiple detection methods catches malware from various angles. A defense-in-depth approach implements security controls that complement each other and provide several layers of protection. For example, most attacks begin with social engineering, specifically phishing. Threat actors then attempt to gain persistence and a foothold, sometimes downloading other malware. From there, they perform various actions such as stealing

information, pivoting to other computers or networks, and even deploying ransomware. A defense-in-depth approach to defend against this could manifest as phishing training for users within your organizations and email filtering for malware and spam. Then, you could deploy an endpoint agent such as AD360 that monitors suspicious and abnormal behaviors. You could deploy extra monitoring tools within your network or anything else that fits your network posture, budget, and business constraints.

However, AD360 has a defense-in-depth implementation within it. When a file arrives on a system, AD360 first compares the hash of the file with known malware or PUP hashes. If the hash isn't within the database, the behavioral engine attempts to classify the file. Following this, AD360 sends the file to the Cloud engine that performs additional analysis and applies detection rules and digital signature checks simultaneously. Finally, if none of these measures classify the file, it moves on to Panda's attestation team, where an attestation analyst analyzes the file. If the file reaches this stage, a classification will always be determined. Whether it's goodware, PUP, or malware, it is up to the analyst.

However, in Q2, attestation analysts were responsible for 9.2% of classifications. On the contrary, the initial signature check was responsible for almost half of all determinations (50.0%). Following this, the subsequent behavioral engine classified 18.2% of all files processed by AD360. The Cloud and defined rules respectively classified 11.1% and 10.1% of files. Finally, a digital signature check invoked a classification between 1% and 2% of the time in Q2 (1.4%). You can observe in the graph the ratio of detections for each technology that shows how AD360 implements a defense-in-depth approach on the endpoint.

Alerts by Technology

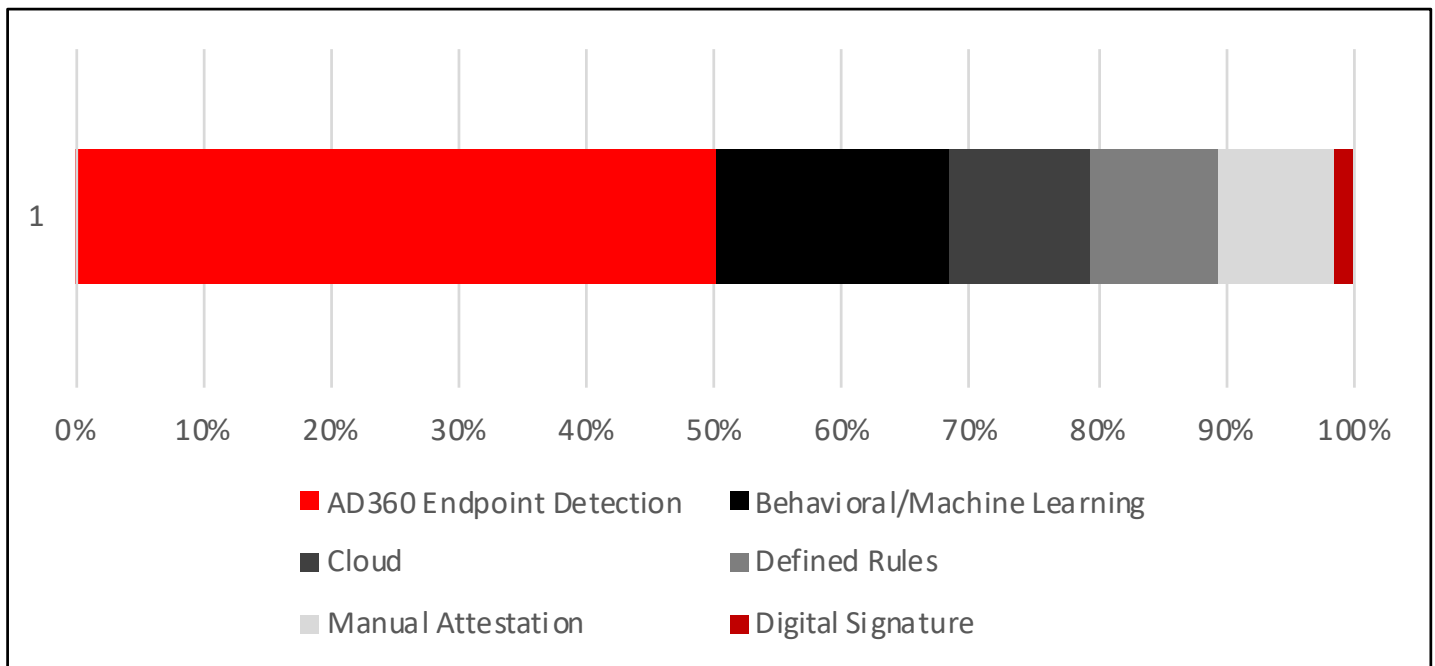


Figure 29. Alerts by Technology

ATTACK VECTORS

Previously, this subsection was titled Exploits. However, we've renamed it to its original name that appeared in 2022 and prior – Attack Vectors. We figured the name suits what it describes, a subset of software threat actors leverage to breach a network. This can come in the form of attacks such as process injection, where malware injects a payload into another trusted process. It could even be a hijacking process or software with a vulnerability. Occasionally, software is trojanized from the source, as we observed with the 3CX installer in late Q1 going into Q2. Read the Top 10 Most Prevalent Malware section for more information about that. Threat actors leverage various tools, tactics, procedures, and attack vectors to achieve their nefarious tasks. The attack vectors that we log are defined below.

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, msixexec.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.

Unlike most other data in the Endpoint section, we've collected





Figure 30. Top Exploited Software

these data points for several consecutive quarters. For the first time, we have omitted the Acrobat attack vector because there weren't any detections to report. This is similar to last quarter when we removed Java because of a lack of detections. Additionally, we no longer track Mozilla detections separately because those detections initially expanded beyond Firefox, but now they are only Firefox again. This leaves us with the five attack vectors: Browsers, Office, Other, Scripts, and Windows.

All of the attack vectors showed a reduction in detections besides the Windows attack vector, which increased by 29.5%. This wasn't the most significant difference from the quarter prior. The Browsers attack vector saw a reduction of 33.2% from the previous quarter, and the biggest surprise was the sharp reduction in detection counts for Scripts at 41.1%, the largest swing of any attack vector. A decrease in Scripts, responsible for most of the attack vector detections, means the overall detections reduced for the quarter. This matches our data for the overall malware frequency above. The other two attack vectors, Office and Other, also saw reduced detection counts at 12.8% and 6.8%, respectively.

Alerts by Exploit Type

The prior Attack Vectors subsection showed how threat actors leverage software to breach, pivot, and gain persistence within networks. As usual, threat actors leveraged scripts the vast majority of the time. This subsection aims to dive deeper and showcase the top exploits attackers use. If the last subsection explained the attack vectors threat actors use to breach networks, then this subsection explains the exact exploits used against these types of software. They're complementary.



Exploit	Alert Count	Description of Exploit
ShellcodeBehavior	16,966	.NET files that allocate and inject payloads directly within the memory of it's own process (Assembly.Load)
NetReflectiveLoader	9,581	Code execution on MEM_PRIVATE pages that do not correspond to a PE
RunPE	3,773	Process Hollowing Techniques
WinlogonInjection	2,334	Remote Code Injection into winlogon.exe process
PsReflectiveLoader1	2,016	Files that leverage PowerShell to allocate and inject payloads directly within the memory of it's own process (E.g. Mimikats) (Local)
ThreadHijacking	926	A process injection technique that allows the execution of arbitrary code in a separate process
ROP1	809	Return Oriented Programming
RemoteAPCInjection	690	Remote code injection via APCs
DumpLsass	418	LSASS Process Memory Dump
HookBypass	413	Detection of memory allocation in base addresses; typical of heap spraying
IE_GodMode	352	GodMode technique in Internet Explorer
DynamicExec	112	Execution of code in pages without execution permissions (32 bits only)
Shellcode_Behavior	69	Code execution on MEM_PRIVATE pages that do not correspond to a PE
APC_Exec	54	Local code execution via APC
JS2DOT	44	.NET Reflective Loading Technique
ReflectiveLoader	22	Reflective executable loading (Metasploit, Cobalt Strike, etc.)
ReverseShell	20	Detection of reverse shell
CVE-2021-26411	7	Microsoft Internet Explorer Memory Corruption Vulnerability

Figure 31. Alerts by Exploit Type

Unfortunately, the results of this subsection are unceremonious. No new exploits made the list, and most of the exploits that remained on the list were indistinguishable from last quarter. For example, RemoteAPCInjection, which is remote code injection via APCs, and JS2DOT, a .NET reflective loading technique, had no change in their rankings. The most significant difference from last quarter was the change in the number of ThreadHijacking detections, which saw this exploit type move up to sixth position, an increase of eight from the quarter prior. The most significant reduction in rankings was from the HookBypass exploit, which is when memory is a dynamic allocation of memory in the base addresses of the file, typically performed via heap spraying. Other than that, the other exploits moved only one or two places. You can view them in the Alerts by Exploit Type table.



RANSOMWARE LANDSCAPE

All of 2022, we observed elevated levels of ransomware detections. At the turn of the year, in Q1 of 2023, we observed a sharp quarter-over-quarter reduction of 73.3% in ransomware detections. In Q4 2022, there were 2,225 ransomware detections; in Q1 2023 there were only 593. Well, that trend continues this quarter, as we observed a ransomware detection count of 465. This is a further 21.6% QoQ reduction and a 72.4% year-over-year difference. Does this mean that ransomware attacks are vastly on the decline? Not exactly. We derive this data from our AD360 licensed customers, a sliver of the overall threat landscape. For example, the following subsection highlights the active double extortion ransomware groups, and we observed a sharp increase in the number of public double extortions.

Extortion Groups

For a few quarters now, we've been tracking and logging the victims posted by ransomware groups that perform double extortions. In some cases, we can uncover victims from public reporting and even from the breached organizations. Some organizations are within jurisdictions that require them to disclose breaches of any kind to local governments. However, most of the data we gather are from these groups' dark web data leak sites (DLS). To be clear, we don't download or analyze the data disclosed by these groups. We only log the victim names, dates, industry sectors the victims operate in, and, in some cases, the extortion amount in US dollars.

These ransomware groups come and go every quarter. Some linger around longer than others, but most don't make it past a year in operations. Q2 saw a record number of new double extortion groups since we began tracking them. Most of these are still operating. Only one of these groups has gone offline or dormant – CrossLock. If we had to highlight one of these groups, it would be MalasLocker. This group appeared with a whopping 172 victims, all using Zimbra servers.

Furthermore, the operators of MalasLocker appear to be hacktivists because their DLS displays a manifesto, discussing how rich people are the root of all of the problems in society. You can view all of the new groups below. We've also injected them into our Public Extortion Groups graph, denoted by an asterisk.

New Groups:

- 8base
- Akira
- BlackSuit
- CrossLock
- CryptNet
- DarkRace
- DungHill Leak
- MalasLocker
- RA Group
- Rancoz
- Rhysida
- Trigona




Figure 32. Ransomware Detections by Quarter

Even though there is a declining trend of ransomware detections for users who use AD360, the number of ransomware groups that publicly extort victims has significantly increased. We observed a 71.8% increase in double extortion attempts by ransomware operators from Q1 to Q2. Which groups are responsible for this increase? Almost all of them.

The groups that had an increase of victims from the quarter prior are:

Group	Q1	Q2	Percentage Increase
BianLian	43	98	127.91%
Bl00dy	0	2	200.00%*
Black Basta	20	76	280.00%
BlackByte	12	19	58.33%
BlackCat	78	137	75.64%
Cuba	0	7	700.00%*
Donut Leaks	0	1	100.00%*
Everest	3	5	66.67%
Karakurt	5	19	280.00%
Medusa Blog	28	38	35.71%
MedusaLocker	0	2	200.00%*
Money Message	6	8	33.33%
Monti	6	12	100.00%
Nokoyawa	0	20	2000.00%*
Play	53	67	26.42%
Qilin	0	24	2400.00%*
Ragnar Locker	2	3	50.00%
Snatch	5	26	420.00%
Vice Society	3	12	300.00%

Figure 33. Increase in Ransomware Groups

*Percentages with an asterisk are displayed for representative purposes. Mathematically, an increase beginning from 0 is undefined.

Notable Ransomware Breaches

Black Basta

ABB – In early May, Switzerland-based technology firm ABB confirmed a network breach. They also confirmed the breach resulted in ransomware and data exfiltration. ABB has over 100,000 employees, around \$30 billion in revenue, and does business internationally for private and public firms, including major departments in the United States such as the Department of Defense (DoD). Later the same month, reporting showed that Black Basta was responsible for the attack.

Rheinmetall – Rheinmetall is a German designer and manufacturer of military armaments. Shortly after the ABB breach above in May, BlackBasta posted Rheinmetall to its dark web data leak site in a double extortion attempt. The defense manufacturer detected the attack in mid-April, allegedly resulting in data encryption and exfiltration.

Viking Coca-Cola – In a busy quarter for BlackBasta, where the group listed 76 victims versus the 20 from the quarter prior, Viking Coca-Cola was listed as one of the victims on the group's dark web data leak site. Coincidentally, BlackBasta listed Viking Coca-Cola on May 17, around the same time as the two victims above. As you may have guessed, Viking Coca-Cola is associated with the Coca-Cola company as one of their largest bottling partners.

BlackByte

City of Augusta, GA – May was a busy month for ransomware operators, as the city of Augusta, Georgia, got caught in the whirlwind of attacks. The city confirmed the attack, but at the time of disclosure, it was uncertain what the cause was. That was until a new group named BlackByte listed them on their dark web data leak site. This attack is notable because this is the second largest city in Georgia, and BlackByte is a new group rumored to demand a ransom of \$50 million. However, the mayor claimed this was false. On the dark web site, BlackByte allowed users to delete the data for \$400,000 and sell it for \$300,000.

BlackCat

Western Digital (W.D.) – If you're familiar with the technology sector, you've likely heard of W.D. They are one of the United States' largest producers of storage systems, such as hard disk drives (HDDs), solid-state drives (SSDs), and other storage-related products. Although the attack happened in late March, much of the data about this attack didn't come to light until April, when the group demanded an "8-figure ransom." What is unique about this attack is that the BlackCat group performed additional blackmail efforts, such as taking a picture of a videoconference of Western Digital's threat-hunting team.

CLOP

MOVEit zero day exploit – Last quarter, we discussed a zero day exploit in GoAnywhere MFT software. This quarter, unfortunately, CLOP is back with another zero-day vulnerability exploitation. This one appears to be much more significant in terms of impact. The group exploited MOVEit software, a secure managed file transfer service similar to GoAnywhere. Two instances are not yet a pattern, but it appears that managed file transfer services are a current target of this group. The MOVEit exploit has left hundreds of orga-

nizations vulnerable by using the software in their organizations and, thus, exposed to the vulnerability. The number of organizations affected is in the hundreds, and it wouldn't be a surprise if the number of known victims moves into the thousands. The number increases by the day.

Play

City of Lowell, MA – The Play ransomware group had an increase of 14 victims from the quarter prior that we know about (Q1: 53; Q2: 67). On the Play dark web data leak site, the group claims to have exfiltrated 5 G.B. of data preceding an encryption event on the city of Lowell, Massachusetts network. The city's Chief Information Officer, Miran Fernandez, led an effort to completely overhaul the city's computer systems, saying that the effort was "the biggest reboot in the city's history." It is uncertain if they paid any ransom.

Rhysida

Government of Martinique – Martinique already dealt with a cyberattack in January of this year that led to the loss of business continuity. A new group named Rhysida struck the government of Martinique again in May, shortly after the group emerged. This led to even more downtime of cyber operations, and Rhysida began leaking data they stole from the networks shortly after that.

Chilean Army (Ejército de Chile) – Around the same time as the attack above, Rhysida breached the Chilean Army with another cyberattack. The weapon of choice was ransomware, and the group also exfiltrated data. They began leaking the data they got from the Army's network. Of course, the data of an army of a nation is never good to have public. The Chilean Army confirmed the attack in early June.

Royal

City of Dallas, TX – The Royal ransomware group, which has ties to the old Conti ransomware group that disbanded, is confirmed to have attacked the city of Dallas, Texas with ransomware. If you were a Dallas resident at the time. In that case, you were possibly affected by this attack as many city services, including emergency services such as ambulances and 911 operations, faced disruption. It is unknown if the city paid the ransom, but we know, via self-admission, that they spent millions of dollars on recovery efforts. So, successful ransomware attacks cost a lot even if you don't pay the ransom.



CONCLUSION & DEFENSE HIGHLIGHTS



CONCLUSION AND DEFENSE HIGHLIGHTS

Even though we stirred the pot a bit this quarter, making it more difficult to compare our threat trends to historical quarters due to our new methodologies, we still learned a lot about the threat landscape. Malware is up quite a bit from a network perspective, but down a little from an endpoint view. Evasive or zero day malware dropped a little too, but seems to be more present in encrypted connections. Ransomware volume is down but the groups are still active, popping new victims every day and perhaps asking for higher ransoms.

Meanwhile, network attacks have dropped a lot – maybe in part due to our new outlier methodology—but a big drop nonetheless. We continue seeing threat actors go after old software vulnerabilities, and CISA seems to confirm that. We also saw cybercriminals increasingly target Linux servers. So what is a concerned IT or security professional to do?

We already gave you some of those tips throughout this report, but let's end with three find tips that can help you against the threats we saw during Q2 2023.

Abandon Abandonware!

During Q2, we saw attackers trying to exploit vulnerabilities in two different old applications that the creators no longer support, both the streaming media server Subsonic and a learning management system (LMS). Obviously, when you can find a free software product that serves a core need for your organization, that is great! That's often the beauty of open source. However, whether it's open or closed source, paid for or not, once a team or company no longer supports a software package, it starts to become a greater risk to you every day, even if it technically still works. If threat actors discover new vulnerabilities, or even old unpatched ones, the project maintainers are not they to update it, leaving you with more and more vulnerable software. Rather than hold onto some tool just because it's easy, free, or convenient, if a software vendor or open-source project abandon a package, you need to immediately begin a plan of how to move out of it, and also make sure to mitigate the risk until you do. For instance, you can use our Firebox to limit network access to this abandonware, to at least prevent Internet-based attackers for threatening it. If you must expose it, use IPS to block any known exploits. But remember, if someone finds a new zero day in old unsupported software, even IPS will not help.

Harden and Protect Linux Servers

During the quarter, we saw four Linux-based threats top our Malware lists, showing that cybercriminals are going after Linux too. Sure, your employees likely don't use Linux often as a desktop, but it makes a very popular server, and when configured that way, may expose network services like SSH and more. Make sure to spend some time hardening your "gold" Linux image, so that it has all unneeded network services disabled. If you are going to open it up to SSH connections, perhaps considering limiting who might access the server with a Firewall access control list (ACL) or even by requiring the access only over VPN. You should also make sure to enforce strong passwords for the server's users, enable MFA on SSH or any remote services, and better yet, enable certificate-based authentication. These simple protections should guard against some of the Linux SSH brute force attacks we've seen.

However, the advice goes for any network service you expose. Look up various hardening guides for every service you open, to make sure you only expose it in the most limited, least privilege way as possible for the use case.





Look at your defenses with a new perspective

If I haven't hammered the perspective theme home yet, this should be the nail in the coffin for it. You should take some time to look at all your cybersecurity defenses – including network and endpoint policies, privileged account lists, exception lists, and more – from a new and updated perspective. I find that for many small, even medium-sized businesses, security often turns into set and forget. Whether it's because of lack of resources or expertise, or other priorities, many small businesses set up policies for various security controls, and if things generally seem to be working, rarely go back to check or adjust them. In doing so, you might forget overly permissive policies you planned on shoring up, or privileged access control lists that have grown to proportions you didn't originally imagine. You probably even will find users, policies, or exceptions you might want to prune based on new knowledge, or changes at your company. Now that you have more insight into what threat actors are doing around the world, take some time to look at your security strategies, and the detailed tactics (policies) you've set in your security controls, to make sure they still apply with all you know today. Finding a great new perspective doesn't really do much unless you act on the knowledge it brings.

We hope you found our Q2 2023 internet threat landscape report interesting, and maybe gleaned a new tip or two. Return next quarter to see how the trends continue or change, and by then we should have some historical context with which to judge any changes. 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 STUIJBERGEN

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.