

GLOBAL THREAT REPORT 2026

1st Edition

“Costly thy habit as thy purse can buy ... For the apparel oft proclaims the man.”

– William Shakespeare, Hamlet

Those lines belong to Polonius, the meddling old counselor in Hamlet, sending his son off to France with a satchel of fatherly advice: buy the best clothes you can afford, because people will judge you by what you wear. For most of history that worked in both directions. A uniform, a badge, a logo on the side of a van told you who you were dealing with before anyone said a word. Our industry was built on the same assumption. We learned what malicious code looked like, wrote its appearance down as a signature, and turned away anything wearing the wrong clothes.

In the first half of 2026, the apparel stopped proclaiming anything. In Q2, 95.72% of the malware we blocked on endpoints appeared on exactly one machine. One victim, one payload, one costume, never worn again. That is not a spray campaign; it is a tailor’s shop. Malware-as-a-service builders and large language models have made a hand-stitched disguise cheaper to produce than an off-the-rack one, and threat actors are ordering one per victim.

Here is the twist, though. All that craftsmanship went into the disguise, not the break-in. The median vulnerability among our 50 most voluminous IPS signatures was first disclosed in 2014, and not one of them targets a flaw from 2025 or 2026. So picture our adversary: a figure in a bespoke suit no one has ever seen before, walking up to your building and trying a key that was cut twelve years ago. Distressingly often, the key still turns.

When it doesn’t turn, he doesn’t force the door. He lifts the keys off your own keyring. Credential access was the most-alerted tactic in our Q1 threat hunting data, led by stealing passwords out of browsers. He then works with the tools already hanging in your utility closet, as Windows living-off-the-land (LotL) binaries kept taking share from malicious scripts. He also arrives in an unmarked vehicle with the windows tinted: 95% of the malware our

Fireboxes caught traveled inside an encrypted connection, and only 20% of deployed Fireboxes inspect encrypted traffic at all.

That is the through line of this report, and it comes with a warning. Volume fell nearly everywhere we measure it, and it would be a serious mistake to read that as calm. The threats that reach you now are fewer but stranger, far more specific to you, and dressed in something you will never see twice. You cannot recognize this adversary by appearance anymore. You must recognize him by behavior, which is where our proactive and machine-learning engines did most of this half’s work.

Welcome to our second biannual Global Threat Report. The mission hasn’t changed: study how attackers behave in the wild, then turn that into defenses you can put to work now. The longer half-year window continues to earn its keep, letting real signal separate itself from the short-lived spikes that used to dominate a single quarter’s view.

In Q2, **95.72%** of the malware we blocked on endpoints appeared on exactly

ONE Machine.

ONE Victim.

ONE Payload.

ONE costume, never worn again.



OUR THREAT REPORT IS BROKEN DOWN INTO THE FOLLOWING SECTIONS:

10 NETWORK-BASED MALWARE TRENDS

This section draws on the Firebox's three network antimalware services, covering everything from the top threats by volume to how much malware slips past legacy defenses. Detection volume fell sharply this half, yet our proactive engine caught more than before – and almost all of it still arrived encrypted.

18 NETWORK ATTACK TRENDS

The Firebox's Intrusion Prevention Service (IPS) blocks known software exploits aimed at servers and clients. Volume collapsed to roughly one blocked attack per Firebox per day, down from nearly five, while the variety of flaws under attack broadened. We also measured how old those flaws really are, watched SQL injection stage a genuine comeback, and followed a generic web shell signature as it became the most widespread network attack on Earth.

25 ENDPOINT THREAT TRENDS

Alongside the Firebox, we inspect the malware blocked on the millions of endpoints protected by WatchGuard's Endpoint Security portfolio and AD360. Total threats held roughly steady while never-before-seen threats kept climbing, nearly every alert touched a single machine, and behavioral engines took over most of the detection work. This section also covers attack vectors, exploit techniques, threat hunting, and a ransomware ecosystem adding new extortion groups every few days.

New clothes call for new eyes. As always, the point of all this data is not to catalog what threat actors did, but to turn it into defenses that match how they now operate. We share tips throughout, and the closing Conclusion and Defense Highlights section pulls out the three strategies that matter most for the way attackers dressed themselves this half.

TABLE OF CONTENTS

TABLE OF CONTENTS

02	Introduction
05	Executive Summary
06	Highlights from H1 2026 Report
07	Network-Based Malware Trends
09	Firebox Feed
10	Malware Trends
12	Top 10 Malware Detections
13	Top 5 Widespread Malware Detections
14	Geographic Threats by Region
16	Individual Malware Sample Analysis
18	Network-Based Attack Trends
19	Top 10 Network Attacks
22	Most-Widespread Network Attacks
24	Network Attack Conclusion
25	Firebox Feed: Defense Learnings
26	Endpoint Threat Trends
28	Malware Frequency
33	Top Malware and PUPs
36	Attack Vectors
46	Threat Hunting
50	Ransomware Landscape
53	Conclusion and Defense Highlights
56	About WatchGuard



EXECUTIVE SUMMARY

The first half (H1) of 2026 is a study in the difference between how much and how many. Nearly every measure of volume in this report went down, and several went down hard. Average network malware detections per Firebox fell 28%. Network attack detections fell roughly 79%, from about five blocked exploits per Firebox per day to about one. Signature-based Gateway AntiVirus hits fell 43%. Read only those numbers and you would conclude the threat landscape took the half off.

Now read the other column. Never-before-seen endpoint malware rose 21.57% from Q4 2025 to Q1 2026 and another 8.01% into Q2, capping a 2065% year-over-year increase. Novel malware grew roughly twenty times year-over-year while total endpoint threats stayed close to flat. The share of endpoint alerts that appeared on exactly one machine climbed from about 88% to 95.72%. Unique IPS signatures firing rose from 381 to 410, the top 10 signatures shrank from 74% to 60% of all detections, and the most targeted 1% of Fireboxes absorbed 46% of attack volume – down from 66%. Proactive IntelligentAV detections rose 49% while its signature-based sibling fell. Attackers may have done less this half but did far more different things.

The explanation runs through the endpoint data. Malware-as-a-service builders and LLM-assisted tooling have made one-off payloads cheap, so threat actors increasingly craft a unique file for each targeted machine rather than blasting one payload at thousands. Our detection mix has moved in step: behavioral and machine-learning engines grew from roughly 40% of all endpoint detections in Q1 to over 67% in Q2. This means about two-thirds of what we stopped in Q2 was caught by AI-related countermeasures rather than anything that recognized the file itself.

The network attack data supplies the counterpoint, and it is a strange one. While the payloads got new, the exploits got old. The median vulnerability under attack in the top 50 signatures was first disclosed in 2014, 31 of the 44 CVE-referenced signatures target flaws at least a decade old, and nothing in the top 50 targets a vulnerability disclosed in 2025 or 2026. All five first-time entrants to our top 50 target flaws are between nine and twenty-three years old, including a twelve-year-old command injection in a discontinued security appliance. In fact, three SQL injection signatures took the Nos. 3 through 5 spots, and Log4Shell still held No. 8 nearly five years after disclosure. Attackers did not get new tools this half. They pointed the old ones at more doors.

Finally, where the old keys fail, adversaries went around our controls rather than through them. Credential access was the top ATT&CK tactic in our Q1 threat hunting data, led by stealing passwords out of browsers, and trust control evasion was the second most-invoked threat hunting rule. Windows LotL binaries kept taking share from malicious scripts as an endpoint attack vector, and process hollowing (RunPE) exploded to 77.1% of all Q2 exploit techniques while AMSI bypass attempts rose roughly fourteenfold. Stolen keys and trusted binaries do not need an exploit at all.



Novel malware grew roughly **twenty times** year-over-year while total endpoint threats stayed close to flat.

HIGHLIGHTS FROM OUR H1 2026 REPORT:

Network-based malware trends: This section draws on the Firebox's three network anti-malware services, covering everything from the top threats by volume to how much malware slips past legacy defenses. Detection volume fell sharply this half, yet our proactive engine caught more than before – and almost all of it still arrived encrypted.

Our “per Firebox” malware results for the various network malware detection services:

- Average total malware detections per Firebox: 903 (28% decrease HoH)
- Average malware detections by GAV per Firebox: 561 (43% decrease HoH)
- Average malware detections by IAV per Firebox: 281 (49% increase HoH)
- Average malware detections by APT Blocker per Firebox: 61 (35% decrease HoH)

95% of malware arrived over encrypted (TLS) connections, but only 20% of Fireboxes inspect it. The TLS share is down just one percentage point from the previous half, and encrypted delivery remains the single largest blind spot in network defense. Gateway AntiVirus detections over TLS are down 52% HoH, while evasive malware over TLS experienced a 38% decline against a clean Q2 2025 baseline.

More than a quarter (28%) of malware evaded signature-based detection. On devices running IntelligentAV and APT Blocker, 28% of malware required proactive techniques to catch. On devices that also inspect HTTPS traffic, that zero-day share rises to 36%, confirming that the most evasive threats disproportionately choose encrypted delivery.

Droppers still own the top of the network malware list. Four of the top 10 families by volume were droppers, led by CMD:Heur.BZC.PZQ.Boxter at 710,699 detections and a related variant, Heur.BZC.PZQ.Boxter.919, at 409,801. The #2 family overall, Application.

Agent.IIQ (605,127 detections), is also a dropper and arrived entirely over encrypted connections. **A third Boxter variant appeared in the encrypted top five.**

Linux, IoT, and proxy abuse featured strongly. A rootkit matching the opensource libprocesshider tool posted 201,808 detections while the IPRoyal Pawns.app residential proxy client posted 184,522. Two Linux coinminer families combined for another 214,801 and a Mirai variant was our most widespread malware detection, peaking at 28.75% of Fireboxes in Sweden.

APAC absorbed half of all per-Firebox network malware. At 50.33%, the APAC region saw roughly double the share of both EMEA (26.24%) and AMER (23.43%).

SQL injection is enjoying a renaissance. Three separate SQL injection signatures occupied positions 3 through 5 and together accounted for over 17% of all network attack detections. Two of them had not appeared in our top 10 since 2024 and early 2025.

The median vulnerability under attack was disclosed in 2014. Twelve years old. Thirty-one of the 44 CVE-referenced in the top-50 signatures target flaws at least a decade old. The oldest reaches back to 1999, and not one targets a vulnerability from 2025 or 2026. Log4Shell held at No. 8 despite an 87% volume decline, and all five first-time top-50 entrants targeted flaws between nine and twenty-three years old. These include Apache Tomcat HTTP PUT (CVE-2017-12615), Apache Struts remote code execution, and a command injection in a discontinued Sophos Web Appliance (CVE-2014-2850).

A generic web shell signature became the most widespread attack in the world. It was detected on 52% of Fireboxes in the Americas and 42% in EMEA, including 75% of reporting Fireboxes in Belgium, 60% in Italy, and 59% in the United States. In Belgium, that saturation came from just 774 detections spread across three-quarters of the country's reporting Fireboxes – indicative of a broad automated sweep rather than a focused assault. Reach and volume diverged more sharply this half than at any point in our records.

New endpoint malware grew roughly twentyfold year over year. New threats per 100,000 active machines rose 21.57% from Q4 2025 to Q1 2026 and another 8.01% into Q2. This compounds last year's spike into a 2065% annual increase. Total endpoint threats, by contrast, declined about 13% in each of two quarters before jumping over 40% in Q2, landing close to where they started.

95.72% of Q2 endpoint alerts appeared on exactly one machine. Up from about 90% in Q1 and roughly 88% in Q4 2025, this is the clearest evidence we have that payloads are being generated per victim rather than distributed en masse.

Two-thirds of Q2 endpoint detections came from AI-based countermeasures. Behavioral and machine-learning detections rose from about 40% of all detections in Q1 to over 67% in Q2, while every other detection technology declined. When every payload is unique, behavior is what is left to recognize.

Windows living-off-the-land binaries kept taking share from scripts. Windows-based vectors rose almost 23% from Q4 2025 to Q1 2026 while scripts fell nearly 20%, driven by a steep drop in PowerShell detections.

Exploit techniques flipped almost completely. Process hollowing (RunPE) surged to 77.1% of all Q2 exploit technique alerts while RemoteAPCInjection, usually our most-observed technique, fell to 5.269%. DumpLsass, ReflectiveLoader, and AmsiBypass each rose roughly tenfold or more, with AMSI bypass attempts up almost fourteenfold.

Credential theft defined Q1; persistence defined Q2. Credential access was the top AT-T&CK tactic in Q1 threat hunting, led by stealing passwords from browsers, with trust control evasion the second most-invoked rule. Q2 shifted toward persistence via registry AutoStart keys and startup folders. Attackers are going around our controls with valid credentials rather than through them.

Supply chain campaigns left fingerprints in the data. Git-related detections spiked in Q1 alongside the TeamPCP GitHub repository compromises. They fell to zero in Q2, when NodeJS detections – driven by the TamperedChef cluster – took their place. Q1 VPN and remote access detections similarly tracked the FortiBleed credential campaign.

Old malware refuses to retire. GuLoader was the most prevalent endpoint malware in Q1, joined by coinminers, TamperedChef, Conficker (2008), and Gh0stRAT, whose source code leaked in 2008. Mimikatz and Floxif appeared in Q2, and AutoKMS licensing-bypass tools made up half of the Q2 top 10 PUPs.

Ransomware is rising again on endpoints, and the extortion ecosystem keeps growing. Ransomware detections on WatchGuard-protected machines rose 11.67% from Q4 2025 to Q1 2026 and about 6% into Q2. However, volumes remain low because these actors are stopped earlier in the kill chain. Public extortion claims dipped 1.39% and then 6.74% but stayed well above historical norms. There were 41 new extortion groups that appeared over the past six months, totaling about one every four to five days, and the top eight groups accounted for 50.79% of the 4,932 extortions we tracked. The top two groups, Qilin and The Gentlemen, accounted for 23.30% of attacks.

That's a high-level summary of the top cyberattack trends we saw during the first half of 2026, but our report includes far more detail – and far more defense advice – that will help you protect yourself from these trends. Read on to learn more.

NETWORK-BASED MALWARE TRENDS



FIREBOX FEED

WHAT IS THE FIREBOX FEED?

The following section of this report is based on threat detections from tens of thousands of WatchGuard Fireboxes deployed around the world that have opted in to sharing the data with us. 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:

- 01 **Gateway AntiVirus (GAV):** Signature-based malware prevention
- 02 **IntelligentAV (IAV):** Advanced AI-based malware prevention
- 03 **APT Blocker:** Sandboxed, behavioral-based malware prevention
- 04 **Intrusion Prevention Service (IPS):** Network-based client and server exploit prevention
- 05 **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



MALWARE TRENDS

Firebox Feed provides anonymized data from Fireboxes around the world. This data from those who have opted into the feed allows us to identify cyberattack trends. We filter this feed and analyze it to identify trends in malware, network attacks, and malicious server activity. Our analysis, along with data from previous periods, provides an overview of threats and recent trending threats. Furthermore, we break the data down by region, and sometimes country, so we can know what to look out for in those areas.

We identify encrypted connections that detect malware what service catches it in the Gateway AntiVirus (GAV), APT Blocker, and Intrusion Prevention Service (IPS) sections. If you only have a few minutes, we provide charts for a quick overview of the threat landscape and details on our analysis. A Firebox configured to provide anonymized feed provides details from the GAV, APT Blocker, and IPS services.

In today's cybersecurity landscape, detecting and mitigating malware threats is essential for ensuring the integrity and security of networks. The Firebox Feed reports provide a critical glimpse into malware detection trends by analyzing proxy details, detection engines, and the pathways through which malware is delivered. This report assesses how malware is detected across various engines, whether the malware traveled through encrypted channels, and the general shifts in malware activity during the first half of 2026. By looking at these detection patterns and the tools used to catch them, we can better understand the evolving strategies malware authors are using to evade detection and how organizations can respond to these emerging threats. If you would like to help us improve this report, we ask that you also enable Firebox feedback.

On average, Fireboxes detected 903 malware hits per device during H1 2026. However, this represents a 28% decrease compared to the previous half-year period. Among specific services, Gateway AntiVirus (GAV) remains a critical defense tool, averaging 561 hits per Firebox. This represented a significant decrease of 43% when compared to the prior period. This decline in signature-based detections occurred alongside a 35% drop in APT Blocker hits (av-

eraging 61 per Firebox), which could suggest reduced activity in certain advanced persistent threats or improved upstream filtering. In contrast, IntelligentAV recorded 281 hits on average, reflecting a strong 49% increase. This growth underscores IAV's growing role in identifying novel and evasive malware strains that signature-based methods miss.

Gateway AntiVirus hits over encrypted channels (TLS) averaged 1,082 detections per Firebox, down 52% from the previous half-year. Evasive malware detections over TLS dropped by 38% if we compare to Q2 2025 Q2 and skip the anomalies number in H2 2025. A whopping 95% of malware threats arrived over TLS-encrypted connections during H1 2026. This extremely high TLS share, combined with the continued (though reduced) presence of evasive threats over encrypted channels, further highlights the importance of HTTPS inspection at the network perimeter.

These trends emphasize the dynamic nature of cyber threats and the necessity for multi-layered defensive strategies. As the data demonstrates, while overall volume declined, proactive engines like IAV captured a larger share of sophisticated threats. Strengthening TLS inspection and maintaining advanced behavioral and machine-learning capabilities remain essential to fortify network security in this ever-changing environment.

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.



IntelligentAV recorded
281 hits on average
reflecting a strong
49% increase.

DETECTION VOLUME DECLINES


	DETECTIONS PER BOX	CHANGE SINCE PREVIOUS HALF-YEAR REPORT
Average total malware hits per Firebox	903	Average detections per Firebox dropped by 28% 
Average Gateway AntiVirus (GAV) service hits per Firebox	561	Basic malware decreased 43% 
Average APT Blocker (APT) hits per Firebox	61	APT Blocker dropped 35% 
Average IntelligentAV (IAV) hits per Firebox	281	IAV hits increased by 49% 
Average GAV w/TLS hits per Firebox	1,082	TLS detections by GAV fell 52% 
Average Evasive malware over TLS hits per Firebox	43	TLS detections of evasive malware dropped by 38%* 
TLS malware %	95.00%	Malware over an encrypted connection decreased one point 

Table 1 Half overview

* Evasive malware detections over TLS averaged 43 hits per Firebox. This represents a 38% decline when compared with Q2 2025. (An anomalous spike in the immediately preceding half-year period would have produced a misleading ~98% drop. We therefore used the cleaner Q2 2025 baseline.)

TOP 10 MALWARE DETECTIONS

The Top 10 Malware Detections table highlights the most frequently detected malware families during the first half of 2026, offering crucial insights into prevailing cyber threats. To ensure the list reflects real-world risks, we complement standard detection analysis with statistical reviews. This process helps filter out detections from controlled scenarios, such as users testing Fireboxes or analyzing malware in safe environments, which do not represent active threats. The result is a curated list of the ten most prevalent malware families encountered in the wild. By closely monitoring these trends, readers can stay informed about emerging threats and adapt their cybersecurity defenses proactively.

In the H1 2026 data, droppers continue to dominate the upper ranks by volume. This pattern is consistent with prior reporting periods where these families act as initial loaders that install further malware on compromised systems. The leading entry, CMD:Heur.BZC.PZQ.Boxter, recorded 710,699 detections and is classified as a dropper. Immediately following is Application.Agent.IIQ with 605,127 detections, also a dropper. We see this again in the encrypted malware section since every detection came from an encrypted connection. A related variant of Heur.BZC.PZQ.Boxter.919, appears next with 409,801 detections. These related Heur.BZC.PZQ.Boxter families suggest sustained campaigns leveraging similar dropper techniques, likely delivering secondary payloads such as information stealers or remote access tools.

Linux-based threats are well represented, reflecting attackers’ continued focus on server, IoT, and opensource environments. Variant.Rootkit.Linux.Agent.7 registers as a hacktool with 201,808 detections. Application.Generic.4048548, identified as Hacktool (IPRoyal), accounts for 184,522 detections. IPRoyal is a residential proxy service whose infrastructure is frequently abused by threat actors to mask malicious activity. Two coinminer families, Variant.Application.Linux.Miner.3 (163,476 detections) and Application.Linux.Generic.28255 (51,325 detections), also appear, underscoring ongoing interest in cryptocurrency mining malware targeting Linux systems.

The remaining entries round out a mixed but familiar set of threats. Trojan.Valyria.8923 contributes another 113,318 detections as a dropper. Heur.Mint.Zard.24, a family that has resurfaced in prior periods, appears as a botnet with 96,602 detections. Finally, Exploit.CVE-2017-11882.Gen (52,782 detections) shows that older Microsoft Office equation-editor exploits continue to find victims.

This distribution shows a clear emphasis on multi-stage infection chains, with droppers serving as gateways to additional threats, alongside specialized Linux tools for reconnaissance, proxying, and resource abuse. We have validated these counts through cross-referencing of detection classifications across engines and normalization for deployment bias among reporting Fireboxes. Several of these families or close variants have appeared in previous periods, though the current volume and specific naming indicate active evolution or renewed distribution in H1 2026.

These top detections by volume set the stage for the remainder of the first half analysis. We will examine the encrypted subset of these threats next, followed by widespread impact and regional patterns.

THREAT NAME	MALWARE CATEGORY	COUNT
Application.Agent.IIQ	Dropper	605,127
Heur.Mint.Zard.24	Botnet	237,327
JS.Acsogenixx.1653.2F5DBE4F	Phishing	35,842
Generic.HTML.Phishing.Q.E4ED0BC2	Phishing	19,706
Heur.BZC.PZQ.Boxter.1041.82E9436C	Dropper	15,100

TOP 5 WIDESPREAD MALWARE DETECTIONS

The Most-Widespread Malware table offers valuable insights into the threats most encountered by Fireboxes, showcasing regional variations and global prevalence. This data reflects the malware strains actively targeting organizations across different geographic zones and provides a clear picture of the current threat landscape. Unlike volume-based rankings, these figures represent the percentage of Fireboxes that detected each family, normalized to account for regional deployment density and reporting.

Notably, Generic.Bash.MiraiB.B0453672 stands out as a Linux-based Mirai botnet variant. It shows the highest concentration in Sweden (28.75%), followed by the Netherlands (16.67%) and Australia (14.21%). This continues the pattern of Mirai-family threats resurfacing and adapting to target IoT and Linux environments across multiple regions. Exploit.MathType-Obfs.Gen also appears prominently, with strong presence in Hong Kong (19.77%), Turkey (16.51%), and Greece (14.86%). Trojan.MSIL.Basic.8.Gen follows with significant activity in Hong Kong, Germany, and Turkey.

JS:Adware.Popunder.14449 demonstrates broader geographic spread, leading in Malaysia, the Dominican Republic, and Brazil. Trojan.Generic.38227153 shows the most concentrated impact in a single country, with India accounting for 40.67% of its detections among the top three.

These widespread threats underscore the importance of strong perimeter protections, timely patching, and user awareness, particularly against botnets, obfuscated exploits, and adware that can serve as entry points for further compromise. Their normalized distribution across Europe, Middle East, and Africa (EMEA) Asia-Pacific (APAC), and Americas (AMER) highlights the need for globally aware but regionally tuned defenses.

MALWARE NAME	TOP 3 COUNTRIES BY %			EMEA %	APAC %	AMER %
Generic.Bash.MiraiB.B0453672	Sweden 28.75%	Netherlands 16.67%	Australia 14.21%	7.09%	8.81%	10.24%
Exploit.MathType-Obfs.Gen	Hong Kong 19.77%	Turkey 16.51%	Greece 14.86%	8.80%	2.27%	3.45%
Trojan.MSIL.Basic.8.Gen	Hong Kong 13.95%	Germany 11.97%	Turkey 11.93%	7.05%	4.02%	2.77%
JS:Adware.Popunder.14449	Malaysia 13.29%	Dominican Re- public 13.24%	Brazil 12.92%	5.12%	2.30%	7.29%
Trojan.Generic.38227153 Heur.BZC.PZQ. Boxter.1041.82E9436C	India 40.67%	Dominican Re- public 14.71%	United Kingdom 8.63%	3.96%	4.94%	6.99%

Figure 1 Most Widespread Malware Detections



GEOGRAPHIC THREATS BY REGION

The Geographic Threats by Region table provides insight into the overall distribution of malware detections across major regions, normalized by the number of reporting Fireboxes in each area. This data highlights how certain regions experience disproportionately higher threat volumes per device, offering critical information for targeted cybersecurity measures and resource allocation. The percentages are weighted to account for differences in Firebox deployment density and reporting participation across regions, ensuring an accurate view of relative threat exposure.

In the first half of 2026, APAC accounts for the largest share at 50.33% of threats per Firebox. This is more than double the share seen in either EMEA (26.24%) or AMER (23.43%). The elevated APAC figure aligns closely with patterns observed in the widespread malware detections, where threats such as Trojan.Generic.38227153 showed extreme concentration in India (40.67% of its top-country impact) and other APAC locations, alongside Mirai variants appearing strongly in Australia. The region also features prominently in several phishing and dropper families from the encrypted and overall top lists.

These normalized regional shares underscore the importance of tailored security strategies. Organizations operating in or supporting APAC environments should prioritize monitoring and controls for the specific families highlighted in prior sections, including Linux/IoT botnets and high-volume droppers. At the same time, the relatively balanced but still substantial shares in EMEA and AMER remind defenders that no region is immune and that globally distributed threats can quickly shift focus.

Zero-Day	28%/72%	Devices with APT Blocker and IAV
Zero-Day with TLS	36%/64%	Devices with APT Blocker and IAV and inspecting HTTPS traffic

CATCHING EVASIVE MALWARE


Previously, we examined where malware is most prevalent and how it is distributed regionally; now, we focus on what types of malware are emerging and how effectively our advanced detection engines identify them. Today’s landscape continues to reveal a significant presence of Zero-Day malware or new and evasive threats that lack identifiable family names due to their unique, ever-changing programming code. While they often employ the same malicious techniques as known malware, their altered code allows them to evade signature-based detection effectively.

In addition to signature-based malware detection via Gateway AntiVirus (GAV), the Firebox uses IntelligentAV (IAV) and APT Blocker. IAV employs advanced file structure analysis and machine learning to detect what GAV misses. APT Blocker analyzes files in a sophisticated sandbox to uncover malicious intent through behavioral observation. These proactive engines are essential for catching polymorphic and never-before-seen threats.

On devices with IAV and APT Blocker enabled, 28% of detected malware evaded signature-based methods. When those same devices also inspect HTTPS traffic, the proportion of zero-day malware rises to 36%. This increase demonstrates that evasive threats are disproportionately delivered over encrypted channels, where they can more easily bypass traditional controls. Fireboxes equipped with IAV and APT Blocker detect substantially more malware overall than those relying on signature-based detection alone. The higher zero-day percentage over TLS further highlights the critical role these advanced engines play in combating sophisticated attacks that leverage encryption to evade detection.

These findings reinforce the necessity of enabling IntelligentAV and APT Blocker across the environment, combined with TLS inspection wherever possible. Without these layered capabilities, organizations risk missing a meaningful portion of the most evasive and damaging threats observed during 2026 H1.

GEOGRAPHIC THREATS BY REGION



INDIVIDUAL MALWARE SAMPLE ANALYSIS

Docu-details1.pdf - OneDrive-Themed Phishing PDF

Our advanced APT detection tool recently blocked a malicious PDF file named Docu-details1.pdf that was being delivered over an encrypted connection. Despite the traffic being protected by HTTPS, the TLS inspection capabilities with APT Blocker identified the threat.

The PDF displays a fake OneDrive-branded image showing the text “KINDLY FOLLOW THE LINK INSTRUCTIONS TO VIEW DOCUMENT” along with the official-looking OneDrive logo. Embedded in the file is a link to the URL [https://cloud-onedriveid-blde23amp.meetliaross\[.\]info/](https://cloud-onedriveid-blde23amp.meetliaross[.]info/), a known phishing domain created to steal Microsoft and OneDrive credentials or deliver malware. The link is currently inactive, though the broader campaign remains active in the wild.

This file attempts to trick recipients by masquerading as a legitimate document shared through OneDrive. It exploits user trust in the Microsoft branding and the common “view document” social-engineering tactic. The successful detection on encrypted traffic demonstrates the value of advanced inspection tools with TLS inspection, as traditional perimeter defenses would likely have missed this threat entirely.

Users should remain cautious with unexpected PDF attachments or emails that claim to originate from OneDrive, SharePoint, or similar cloud services. Always verify the destination of any link before clicking, enable advanced threat protection and document sandboxing where available, and report any suspicious files immediately without opening them. Our threat intelligence team continues to monitor this campaign closely.

The PDF displays a fake OneDrive-branded image showing the text “KINDLY FOLLOW THE LINK INSTRUCTIONS TO VIEW DOCUMENT” a known phishing domain created to steal Microsoft and OneDrive credentials or deliver malware.




Figure 2 docuphish

Variant.Rootkit.Linux.Agent.7

The detection Gen:Variant.Rootkit.Linux.Agent.7, matches the compiled form of the open-source lib-processhider proof-of-concept. This user-mode rootkit hides specific Linux processes from standard administrative tools, allowing follow-on malware such as cryptominers or reverse shells to remain invisible while still consuming CPU and maintaining network connections.

The technique abuses the Linux dynamic linker through LD_PRELOAD. An attacker with root privileges places the shared library (commonly named something like libprocesshider.so or libnetresolv.so) into a system path and appends its full path to /etc/ld.so.preload. Every subsequently launched process that links against libc then loads the malicious library first. The library overrides the readdir() function. When tools such as ps, top, htop, or lsof enumerate the /proc filesystem, the hooked readdir() examines each directory entry. If the process name matches a hardcoded target (for example “xmrig” or a custom payload name), that PID directory is silently dropped from the results. The process continues running, yet it never appears in the output of conventional utilities.

This approach requires no kernel modules or binary patching, making it lightweight and effective. Real-world campaigns, including certain cryptomining operations and earlier TeamTNT activity, have reused nearly identical code to conceal mining processes after initial compromise. TeamTNT is a financially motivated cybercrime group, believed to be German speaking based on language artifacts in their tools and scripts.

The last time we observed similar process-hiding libraries, they arrived primarily over encrypted channels and targeted both cloud and on-premises Linux servers.

From a defender standpoint, reliance on /proc-based monitoring alone is insufficient. Organizations should enable continuous inspection of /etc/ld.so.preload for unexpected entries, deploy tools that examine system calls or raw kernel data like Watchguard Endpoint Security, and ensure TLS inspection is active so that the droppers delivering this rootkit cannot slip past network defenses. Integrity checks on system libraries and preload configurations close the remaining window before the rootkit can fully embed.

Application.Generic.4048548

Application.Generic.4048548 is a generic detection for an ARM-compiled update of the Pawns.app client from IPRoyal. The application combines surveys and offers with a bandwidth-sharing feature that turns the device into a residential proxy node, routing third-party traffic through the user’s connection for small payments.




Figure 3 pawns.app.dropper

On ARM devices it installs persistently, runs a background service, and forwards external traffic. While not a traditional trojan, security products flag it as a PUP or riskware due to proxy behavior. The app is frequently bundled with other software, so its presence can indicate that additional unwanted or malicious programs have also been installed on the device.

In a corporate environment the risks are significant. The device becomes an uncontrolled proxy exit node, allowing third-party activity such as scraping or fraud to originate from the organization’s IP addresses. This can damage IP reputation, trigger blacklisting, consume bandwidth, and violate acceptable-use policies. Encrypted proxy traffic also complicates monitoring.

Treat detections of Application.Generic.4048548 as unauthorized software. Block related domains and processes, remove the application, and investigate the device for additional malware that may have been delivered alongside it.

NETWORK-BASED ATTACK TRENDS

NEW NETWORK ATTACKS

The Intrusion Prevention Service (IPS) on Firebox appliances detects and blocks exploit attempts against vulnerabilities in network-connected applications. This includes both client applications, like web browsers, and server applications, like websites and self-hosted tools. In the first half of 2026, Fireboxes with IPS enabled reported a steep decline in network attack volume. Each Firebox saw an average of 181 detections during the period, down roughly 79% from the 845 detections per Firebox in the second half of 2025. Put another way, the average Firebox blocked about one network attack per day this period, down from nearly five per day in the previous half.

The decline was top-heavy. The signatures that dominated late 2025 collapsed the hardest: detections for the dotCMS access control weakness that had led our list since mid-2025 fell 89%, and the HAProxy header-bypass exploit that ranked #2 last period fell over 99%, dropping from more than 322,000 detections to fewer than 2,500. Meanwhile, the long tail broadened. We saw 410 unique signatures fire during the half, up from 381, and the top 10 signatures accounted for 60% of all detections, down from 74%. The concentration of attacks across networks eased in the same way: the most-targeted 1% of Fireboxes absorbed 46% of all detections, down from 66% in H2 2025. In short, the mass-scanning campaigns that generate towering detection counts against a handful of signatures went quiet, while attackers probed a wider variety of vulnerabilities at lower intensity. As always, common web application flaws such as directory traversal, SQL injection, and cross-site scripting made up the overwhelming majority of what we blocked.

TOP 10 NETWORK ATTACKS

SIGNATURE	TYPE	NAME	AFFECTED OS	PERCENTAGE
1059877	Exploits	WEB Directory Traversal -8	Windows, Linux, FreeBSD, Solaris, Other Unix	13.40%
1133539	Web threats	WEB dotCMS CMSFilter assets Access Control Weakness (CVE-2020-6754)	Network Device, Others	10.96%
1058077	Web threats	WEB SQL injection attempt -2.u	Windows, Linux, FreeBSD, Solaris, Other Unix, Mac OS	8.09%
1058468	Web threats	WEB SQL injection attempt -1.b	Windows, Linux, FreeBSD, Solaris, Other Unix, Mac OS	4.84%
1054837	Web threats	WEB SQL injection attempt -25.u	Windows, Linux, FreeBSD, Solaris, Other Unix	4.71%
1055396	Web threats	WEB Remote File Inclusion /etc/passwd	Windows, Linux, FreeBSD, Solaris, Other Unix	4.26%
1230275	Web threats	WEB Cross-site Scripting -9	Windows, Linux, FreeBSD, Solaris, Other Unix, Network Device	3.94%
1059958	Web threats	WEB Apache log4j Remote Code Execution -2.h (CVE-2021-44228)	Windows, Linux, FreeBSD, Other Unix	3.72%
1132896	Web threats	WEB Directory Traversal -27.u	Windows, Linux, Others	3.59%
	Exploits	WEB Remote Shell Command Execution -1	Linux, FreeBSD, Solaris, Other Unix	2.84%

A generic directory traversal signature (1059877) took the #1 spot this period, reclaiming the position it previously held in Q4 2024 and Q1 2025. Its own detection count roughly halved compared to H2 2025. It reached the top not by growing, but by declining more slowly than everything that previously sat above it. The dotCMS CMSFilter weakness (CVE-2020-6754) surrendered the #1 position it had held since mid-2025, though at nearly 11% of all detections it remains firmly entrenched at #2. The most notable shift in the top 10 is the return of SQL injection in force. Three separate SQL injection signatures occupy the #3 through #5 positions. Signature 1133539 returned to the top 10 for the first time since Q1 2025, and signature 1058077 came back after last appearing in Q3 2024. Combined, the three SQLi signatures accounted for over 17% of all detections, evidence that the oldest trick in the web attack playbook is enjoying a renaissance. SQL injection has appeared in every OWASP Top 10 since the list began, and attackers clearly still find enough unsanitized input fields to make automated SQLi campaigns worthwhile.

Log4Shell (CVE-2021-44228) refuses to leave. Nearly five years after disclosure, it held the #8 position, though its absolute volume fell 87% from H2 2025. The persistence of Log4j exploitation long after global patching campaigns illustrates how long vulnerable systems linger on the internet once a library is embedded deep in software supply chains.

Finally, signature 1132896 (WEB Remote Shell Command Execution -1) cracked the top 10 by volume for the first time. Readers of our H2 2025 report will recognize it: it debuted in our top 50 just last period and immediately became one of the most widespread attacks we track. Its climb continues this period. More on that in the most-widespread section below, where it now sits at #1.

Two former heavyweights fell dramatically. The SHELLCODE NOP Sled signature dropped from #3 to #24, and the HAProxy header bypass (CVE-2023-25725) fell from #2 to #36 by volume. Remarkably, HAProxy remains in the top 5 most-widespread list despite that collapse. Attackers appear to have shifted from flooding a small number of targets with HAProxy exploit attempts to probing broadly at low intensity.

NEW DETECTIONS IN THE TOP 50

Five IPS signatures appeared in our top 50 by volume for the first time ever this period. True to form for network attacks, none of them are new vulnerabilities. Every one targets a flaw that is between nine and twenty-three years old.

SIGNATURE	TYPE	NAME	AFFECTED OS	RANK
1054796	Web threats	WEB HTTP Host Header Buffer Overflow	Windows, Linux, FreeBSD, Solaris, Other Unix, Mac OS	19
1134131	Web threats	WEB Apache Tomcat HTTP PUT Windows Remote Code Execution -2 (CVE-2017-12615)	Windows	27
1133533	Web threats	WEB Apache Struts Dynamic Method Invocation Remote Code Execution -1.b	Windows, Linux, FreeBSD, Other Unix, Mac OS	37
1055091	Web threats	WEB HTTP Directory Traversal -9	Windows, Other Unix	45
1112626	Web threats	WEB Sophos Web Appliance Sophos-Config Write Command Execution -2 (CVE-2014-2850)	Windows, Linux, Other Unix	47

Signature 1134131: Apache Tomcat HTTP PUT (CVE-2017-12615)

Back in 2017, researchers disclosed that Apache Tomcat servers on Windows with the HTTP PUT method enabled (via the read-only initialization parameter set to false) would accept an uploaded JSP file and happily execute it: a one-request path from unauthenticated visitor to remote code execution. Nine years later, attackers are still scanning for misconfigured Tomcat instances. Exploitation is trivial and fully automated in common attack frameworks, making this a staple of opportunistic scanning against any organization self-hosting Java applications.

Signature 1133533: Apache Struts Dynamic Method Invocation RCE

This signature catches exploit attempts against a family of remote code execution vulnerabilities in the Apache Struts 2 framework from 2016 (CVE-2016-3081 and related), with detection logic that also covers the infamous CVE-2017-5638, the Jakarta Multipart parser flaw exploited in the 2017 Equifax

breach. Struts vulnerabilities have caused some of the most damaging breaches on record, and the appearance of this signature in our top 50 for the first time suggests attackers are once again sweeping the internet for unpatched Struts applications, most likely legacy line-of-business apps that were never migrated.

Signature 1112626: Sophos Web Appliance Command Execution (CVE-2014-2850) Perhaps the most interesting new entry targets a security product. CVE-2014-2850 is a command injection vulnerability in the administration interface of the Sophos Web Appliance, allowing an attacker to execute arbitrary commands on the appliance itself. Security appliances and network edge devices have become prime targets in recent years. They sit at privileged positions on the network, are frequently exposed to the internet, and are often left running long after their vendors end support. A twelve-year-old flaw in a discontinued security appliance appearing in our top 50 says everything about how long unsupported edge devices survive in production.

The remaining two new entries, 1054796 and 1055091, are generic signatures for HTTP Host header buffer overflows and directory traversal attempts, respectively, with CVE references dating back as far as 2003. Generic signatures like these catch broad classes of malformed requests rather than a single product flaw, and their arrival in the top 50 is consistent with the wider diversity of low-volume probing we saw this period.



HOW OLD ARE THE VULNERABILITIES UNDER ATTACK?

The five new arrivals hint at a broader pattern, so this period we measured it. Of the 50 top signatures by volume, 44 reference at least one specific CVE (the remaining six are generic detections for classes of malformed traffic). For each of those 44, we took the year of the earliest CVE it covers and asked a simple question: how old is the flaw that attackers are actually targeting?

The median vulnerability under attack in the first half of 2026 was first disclosed in 2014. That’s a twelve-year-old flaw. Thirty-one of the 44 CVE-referenced signatures target vulnerabilities at least a decade old, and the oldest references reach all the way back to 1999. At the other end of the timeline, the newest CVE referenced anywhere in the top 50 dates to 2024. Not a single signature in the top 50 targets a vulnerability disclosed in 2025 or 2026.

That last point deserves emphasis. It doesn’t mean new vulnerabilities aren’t being exploited; high-profile zero days tend to be used quietly and narrowly before signatures and patches exist. But at internet scale, the attacks that actually reach the average network are overwhelmingly recycled. Exploit code for a decade-old flaw is free, stable, and bundled into every scanning framework, and enough unpatched targets remain to keep those campaigns profitable. For defenders prioritizing patch backlogs, this chart is the argument: the boring, ancient CVEs are the ones doing the volume.

Disclosure era of earliest CVE	Signatures in top 50
1999-2004	5
2005-2009	3
2010-2014	18
2015-2019	10
2020-2026	8

MOST-WIDESPREAD NETWORK ATTACKS

SIGNATURE	Name	TOP 3 COUNTRIES BY %			AMER %	EMEA %	APAC %
1132896	WEB Remote Shell Command Execution -1	Belgium 75.00	Italy 59.92	USA 58.70	52.13	41.56	18.24
1136822	WEB dotCMS CMSFilter assets Access Control Weakness (CVE-2020-6754)	Germany 48.60	Brazil 32.17	Spain 12.50	11.00	27.03	8.81
1059877	WEB Directory Traversal -8	Germany 36.21	Italy 18.22	Belgium 14.06	9.12	22.53	22.64
1132381	WEB-CLIENT Javascript Obfuscation in Exploit Kits - 44 (Possible Exploit Kit)	USA 34.93	United Kingdom 22.58	Italy 14.98	26.96	12.71	8.81
1231780	WEB HAProxy h1_headers_to_hdr_list Empty Header Name Access Control Bypass (CVE-2023-25725)	United Kingdom 26.67	Canada 25.18	USA 20.72	20.12	14.39	21.38

The most-widespread list contains no first-time entries this period: all five signatures return from H2 2025. What changed is the order, and the story at the top.

The web shell signature (1132896) completed a remarkable trajectory. It first entered our top rankings in H2 2025, appearing in both the top 50 by volume and the most-widespread list simultaneously. Six months later it is the single most widespread network attack in the world, detected on 52% of Fireboxes in the Americas and 42% in EMEA. The country-level numbers are striking. Three out of every four reporting Fireboxes in Belgium saw at least one detection, along with 60% in Italy and 59% in the United States. Because this is a generic signature that matches common web shell command strings rather than one specific exploit, its ubiquity tells us that web shell deployment attempts are now part of the standard playbook for automated attacks worldwide. Web shells give attackers persistent access after an initial exploit, so a blocked detection here often means an earlier exploit attempt got far enough to try to establish a foothold.

The exploit kit JavaScript obfuscation signature (1132381) climbed from #5 to #4 on the widespread list while doing something we rarely see: it fell out of the top 50 by volume entirely. A signature that doesn't even rank among our 50 most voluminous detections still touched enough networks worldwide to be the fourth most ubiquitous attack on Earth, including 35% of Fireboxes in the US. Its profile is unchanged from last period: broad, low-frequency detections consistent with drive-by and malvertising campaigns that touch many networks without hammering any single one.

The dotCMS and directory traversal signatures continue to be a European story, led by Germany at 49% and 36% of Fireboxes respectively. And HAProxy rounds out the list, the same signature that fell off a cliff by volume, dropping from #2 to #36 while slipping one spot to #5 for reach.

INTENSITY VS. REACH

Volume rankings and widespread rankings measure two different things, and this period they diverged more sharply than at any point in our records. The chart below plots each of the five most-widespread attacks by its rank on both lists.

SIGNATURE	NAME (SHORT)	VOLUME RANK	WIDESPREAD RANK
1059877	Directory Traversal -8	1	3
1136822	dotCMS CMSFilter	2	2
1132896	Web Shell Command Execution	10	1
1231780	HAProxy Header Bypass	36	5
1132381	Exploit Kit JS Obfuscation	Outside top 50	4

Directory Traversal -8 and dotCMS sit where volume and reach agree: they hammer many networks hard. The other three tell the more interesting story. The web shell signature reaches more networks than anything else while ranking only tenth by count. HAProxy and the exploit kit signature barely register by volume yet blanket the globe. Campaigns like these are easy to overlook in a SOC dashboard sorted by count, which is precisely why we track both dimensions. A signature with enormous reach and tiny per-network volume is the fingerprint of an automated campaign knocking once on every door, and one blocked knock on your network means the same knock happened on thousands of others.

COUNTRY SPOTLIGHT: BELGIUM

Belgium appeared in none of our top three country lists in H2 2025. Six months later, it tops the most notable one: 75% of all reporting Fireboxes in Belgium saw the web shell signature (1132896) at least once, the highest saturation we recorded for any signature in any country this period.

Here is the twist: all of that reach came from just 774 detections. Spread across three-quarters of Belgium’s reporting Fireboxes, that is a handful of blocked attempts per affected network, the signature of a broad automated sweep rather than a focused assault.

Belgium’s volume story is the exact opposite. The country ranks fifth world-wide in raw top 50 attack volume with 18,459 country-attributed detections, and 71% of that total comes from a single SQL injection signature (1132793, WEB SQL injection select from attempt -5.h). In fact, Belgian networks absorbed 69% of that signature’s entire global volume. One concentrated SQLi campaign against Belgian targets, plus one worldwide web shell sweep that happened to reach Belgium most thoroughly, put a country of its size fifth on the global volume chart and first on the saturation chart at the same time.

The lesson for defenders extends well beyond Belgium: reach without volume means everyone is being probed, and volume without reach means someone specific is being hunted. Belgium spent the half experiencing both at once.

75%

of all reporting Fireboxes in Belgium saw the web shell signature (1132896) at least once.




NETWORK ATTACKS BY REGION

The regional picture inverted this period. In H2 2025, the Americas accounted for just over half of the per-Firebox detection weight; in H1 2026, they fell to last place at 26%, while APAC surged from 21% to the top spot at 38% and EMEA climbed to 36%. For APAC, this is a return to form rather than a first: the region also led this metric in three quarters across 2024 and early 2025, making H2 2025’s Americas-led half look like the recent outlier. On raw volume, the shift is even more lopsided. EMEA generated 75% of all detections (523,272 of 697,088), with the Americas at 20% and APAC at under 5%.

The two views tell one consistent story. EMEA’s dominance in raw volume comes from having the most reporting devices (2,690 of 3,858 IPS-enabled Fireboxes) combined with heavy campaign activity, and the country-level data bears this out: Germany alone accounted for 42% of country-attributed top 50 volume, with Italy second at 14%. APAC’s per-Firebox lead comes despite the region’s small footprint: its 159 reporting Fireboxes each averaged more detections than devices in any other region, suggesting intense scanning pressure against the region’s networks. We would caution that, with a small population of reporting devices, a handful of heavily targeted networks can move the APAC average substantially. The Americas’ decline stands out. After topping this metric in H2 2025, per-Firebox attack pressure in the region fell to a distant third.

NETWORK ATTACKS BY REGION

Average % IPS Detections per Firebox



CONCLUSION

The first half of 2026 was quieter in volume but broader in geographic reach. Attack volume fell to about one blocked detection per Firebox per day, down from nearly five, yet the diversity of exploited vulnerabilities grew, five signatures entered our top 50 for the first time, and the most widespread attack in the world, a generic web shell detection, reached three-quarters of reporting networks in Belgium. The median vulnerability under attack was disclosed in 2014, and nothing in the top 50 targets a flaw from 2025 or 2026. Attackers did not get new tools this period; they pointed the old ones at more doors.

For defenders, the lessons are unchanged but sharpened. Patch the boring, ancient flaws first. Directory traversal, SQL injection, and long-disclosed RCEs in Struts, Tomcat, and edge appliances are still what most often affect networks. Retire unsupported edge and security appliances, which attackers now treat as front doors. Watch reach as well as volume, because the campaigns that touch every network rarely top the count-sorted dashboards. And keep layered protections like IPS enabled everywhere: a period of lower volume is not a period of lower risk when the probing is this widespread.



FIREBOX FEED

FIREBOX FEED DEFENSIVE TIPS

Lower volume is not lower risk. Attack volume fell at the perimeter in the first half of 2026, but what remained was harder to catch and older than ever. Malware hid almost entirely inside encrypted connections, evasive samples grew as a share of detections, and the exploits reaching real networks overwhelmingly targeted vulnerabilities discovered a decade or more ago. The following practices address the specific gaps attackers leaned on this period.

01. BRING LINUX SERVERS AND DEVICES UNDER REAL MONITORING

OBSERVED: Linux threats claimed an unusual share of this period's top malware: a user-mode rootkit built from the open-source libprocesshider tool ranked fourth with over 200,000 detections, joined by two Linux coin-miners and an ARM proxyware client that silently turns devices into residential proxy exit nodes. A Mirai botnet variant was the most widespread malware family in the world, and the most widespread network attack, a web shell command execution signature, targets Linux and Unix systems. Attackers clearly consider Linux servers, IoT devices, and network appliances soft targets: high-value, always on, and rarely watched as closely as Windows endpoints.

ACTION: Extend the same defensive rigor to Linux that Windows fleets already get. Deploy endpoint protection that examines system calls and kernel data rather than relying on /proc-based tools, which the rootkits we caught are specifically built to deceive. Segment IoT and embedded devices away from production networks, change default credentials, and treat unexpected outbound proxy traffic or CPU spikes on Linux hosts as investigation triggers, since miners and proxyware monetize exactly the machines nobody is looking at.

02. ATTACK THE DECADE-OLD PATCH BACKLOG AND RETIRE END-OF-LIFE DEVICES

OBSERVED: The median vulnerability targeted by our top 50 network attacks was disclosed in 2014, and not one targets a flaw from 2025 or 2026. New arrivals in the top 50 included a 2017 Tomcat flaw, 2016 Struts RCEs, and a 2014 command injection in a discontinued security appliance, while a 2017 Office equation-editor exploit still ranks in the top 10 malware. Attackers are not innovating; they are sweeping for what was never fixed.

Action: Prioritize patching by exploitation evidence, not CVSS score or disclosure date. The boring, ancient CVEs in web frameworks, application servers, and Office documents are what reach networks. Inventory internet-facing legacy applications and retire unsupported edge and security appliances entirely. A device the vendor no longer patches is a permanent open door that no amount of monitoring compensates for.

03. HUNT FOR WEB SHELLS AND HARDEN AGAINST THE INJECTION BASICS

OBSERVED: The most widespread network attack in the world this period was a generic web shell detection. It reached three of every four reporting networks in Belgium and over half in the US and Italy – typically at just a handful of attempts per network. Web shell deployment is now part of the standard automated playbook, and a blocked attempt often means an earlier exploit got far enough to try planting one. Meanwhile, three SQL injection signatures together drove over 17% of all detections.

ACTION: Assume your web servers are being probed for shell deployment regardless of size or industry. Disable unused HTTP methods like PUT, apply file integrity monitoring to web-accessible directories, and alert on new executable or script files appearing in web roots. Close the front doors that lead there: parameterized queries and input validation against SQL injection, and IPS enabled on every segment hosting web applications, since reach-based campaigns knock once on every door and count on some being unwatched.

ENDPOINT THREAT TRENDS

WatchGuard collects and aggregates endpoint data from a myriad of sources. Typically, we intentionally exclude WatchGuard product information and focus on the data. But understanding the endpoint ecosystem changes in the first half of 2026 will provide some much-needed context for the rest of this section.

In January 2026, WatchGuard expanded on its MDR service with Open MDR (Managed Detection and Response), placing the Firebox, AuthPoint, EPDR (Endpoint Protection, Detection, and Remediation), and any other third-party tools in one centralized location for easy management. This increases the scope of data ingestion and protection and allows MSPs (Managed Service Providers) to deliver managed security across the entire environment.

The EPDR portfolio also received a refresh at the beginning of Q2, placing it into four tiers: Endpoint Security Basic, Prime, 360, or Elite. Endpoint Security Basic was formally Endpoint Protection Platform (EPP); Endpoint Security Prime is a new tier tucked in between what was formally EPP AND EPDR; EPDR is now Endpoint Security 360; and Endpoint Security Elite. You can find more information about these tiers on the WatchGuard website, but the most important thing to know with these tier changes with respect to this report is that each tier includes AI-powered Endpoint Detection and Response (EDR) by default.

The additions to WatchGuard’s MDR service and the expansion of the EPDR into the new Endpoint Security portfolio, including AI by default, means that there’s more coverage on WatchGuard protected endpoints. A large chunk of that coverage includes automated or semi-automated detection and response. This is reflected in the data.

Most of the data coverage for this biannual report is similar to the previous one. However, the last report didn’t include threat hunting data and it has returned for this iteration. Additionally, we’ve expanded on the threat hunting data by including the top 25 threat hunting rule invocations instead of the top 10. The threat hunting MITRE ATT&CK matrix mapping has returned too.

DATA ABSTRACT

- Total malware threats
- New malware threats per 100k active machines
- The number of alerts by the number of machines affected
- The number of alerts by which WatchGuard technology invoked the alert
- The top 30 affected countries each quarter
- The top 10 most-prevalent malware
- The top 10 most-prevalent Potentially Unwanted Programs (PUPs)
- Attack vectors
- Cryptominer detections
- Alerts by exploit type
- Top 25 threat hunting rule invocations
- Threat hunting MITRE ATT&CK tactics and techniques
- Ransomware detections
- Ransomware double extortion landscape

MALWARE FREQUENCY


Although the report has gone from quarterly to biannually, most of the data within the endpoint section is still defined in terms of quarters. It provides a good balance between too much granular data (monthly) and not enough data to uncover trends (biannually). Quarters also tend to follow the seasons, which affects how people behave. For example, total threats tend to decrease in Q2 because many people take time off from work. Yes, even hackers take vacations.

Malware Frequency is split into two parts: Total Malware Threats and New Threats. Total Malware Threats is the total number of malware hashes—a unique payload— that was blocked. If that exact payload is blocked on another system, it doesn't iterate the Total Malware Threats counter because it's the same payload. On the other hand, if a polymorphic worm that alters its code subtlety with each unique system it touches is blocked on all systems, each unique iteration of that worm is considered a new threat because it has a completely different file hash value.

New Threats is a subset of Total Malware Threats. It's the sum of all hashes that we've never seen before. We skew this number to be represented in terms of a typical large organization (100,000 active machines) because it's more easily digestible. Every other data point is also a derivative of total malware threats, and thus, drives much of what is seen in the rest of the section.

In the previous report for the second half of 2025, we showed new threats from Q3 to Q4 were up substantially, increasing 1548.48%. We were eager to find out if this was a one-off quarter or the beginning of a new trend, and the data is clear. Not only were new threats up from Q4 2025 to Q1 2026, but they still increased at a moderate pace into Q2 2026. From Q4 2025 to Q1 2026, new threats increased 21.57% and increased another 8.01% from Q1 to Q2. This is after a previous spike, compounding from last year. Overall, we've observed roughly 20 times more newer malware threats (2065% increase) year over year.

H2 2025 - H1 2026 QoQ New Malware Threats Per 100k Active Machines Table




NEW THREATS BLOCKED

Q3 99
Q4 1632
Q1 1984
Q2 2143

Considering new threats is a subset of the total malware threats, it would be reasonable to assume that if the number of newer threats has increased substantially, then the total number of threats would increase in tandem. However, that's not the case, and that tells a story of the cybersecurity landscape. We're continuously observing unique malware payloads that are increasingly created for a specific victim machine. These payloads are led by the widespread use of malware-as-a-service (MaaS) ecosystems, which include builders that create ad-hoc payloads. Each of these produces a unique payload, a unique threat. Additionally, payloads are crafted specifically for each victim after some type of persistence has been created. In other words, malware isn't spammed or brute forced as often as in the past and the data shows it. The Top 10 Malware section later in the report unveils some of these MaaS families responsible for much of the mayhem too.

The total malware threats steadily decreased by about 13% each quarter in 2H 2025 (13.70% from Q3 to Q4 and 13.02% from Q4 to Q1) before suddenly reversing course and jumping by more than 40% from Q1 to Q2 2026. Since the previous quarters had steadily declined, the total number of threats is similar to Q3 2025 levels despite the number of new threats increasing. All in all, we're seeing around the same number of overall threats for the past year, but a significantly larger number of these threats are new.

H2 2025 - H1 2026 QoQ Total Malware Threats Graph



TOTAL MALWARE THREATS

Q3 29,348
Q4 25,326
Q1 22,028
Q2 30,893

DEFINED SCHEMA

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




ALERTS BY NUMBER OF MACHINES AFFECTED

The data in the Alerts by Number of Machines Affected subsection contains the total malware threat data filtered by how many machines the payload appeared on. For example, if a specific threat appeared on only one machine throughout the quarter(s), it’s added to the 1 Number of Machines data bucket. If it appears on 15 machines, it’s placed into the >= 10 & < 50 bucket. The schema for how this data is filtered is in the call-out.

The total number of threats being relatively stagnant for the past year, coupled with an increasing number of new threats, means that malware payloads are being created on the fly for each suspected victim. Therefore, if each payload is truly unique, the number of machines a payload would appear on should be one – the victim’s machine. This is corroborated with the filtered data. From Q4 2025 to Q1 2026, the number of threats appearing on one machine remains at approximately 88% with a shift closer to the 90% mark. That threshold was quickly surpassed from Q1 to Q2, when the number of threats appearing on one machine increased to 95.72% of all alert composition.

Q4 2025 TO Q1 2026 ALERTS BY NUMBER OF MACHINES AFFECTED TABLE

NUMBER OF MACHINES	Q4 ALERT COMP.	Q1 ALERT COMP	% DIFF FROM Q4
1	88.19%	88.99%	0.81%
>=2 & <5	9.18%	8.04%	-1.14%
>=5 & <10	1.59%	1.86%	0.28%
>=10 & <50	0.89%	0.95%	0.06%
>=50 & <100	0.09%	0.09%	0.01%
>=100	0.08%	0.07%	-0.01%




ALERT COMPOSITION — Q1 VS. Q2

NUMBER OF MACHINES	Q1 ALERT COMP.	Q2 ALERT COMP.	% DIFF FROM Q1
1	88.99%	95.72%	+6.73%
>= 2 & < 5	8.04%	3.15%	-4.89%
>= 5 & < 10	1.86%	0.64%	-1.23%
>= 10 & < 50	0.95%	0.43%	-0.51%
>= 50 & < 100	0.09%	0.04%	-0.06%
>=100	0.07%	0.03%	-0.04%

Q4 2025 to Q1 2026 Alerts by Number of Machines Affected Table

Q4 2025 to Q1 2026 Alerts by Number of Machines Affected Graph



The primary takeaway from the first few sections sothus far is that threat actors are decreasingly using spammy tactics and are increasingly crafting payloads specifically for to a given network environment and endpoint. The increasing rise of MaaS ecosystems and the addition of expedited automated scripting using large language models (LLMs) with crafty prompt engineering has resulted in malware authors being able to create tools and payloads to assist in breaches. AI LLM’s, as of present, has-tened what is commonly referred to as the cat-and-mouse game in cybersecurity. Threat actors create novel tools and techniques, and cybersecurity practitioners create countermeasures in delayed tan-dem to remediate the gaps. Only now, there’s more mice and more cats.

DEFENSE IN DEPTH

Defense in Depth is where much of the prior AI and machine learning foreshadowing comes to frui-tion. Instead of filtering the total malware threats by machine occurrence, we collect what technology solution blocked the threat. There are six primary mechanisms that prevent threats on WatchGuard protected endpoints, and each functions differently. Effectively creating a defense in depth posture on each endpoint. Those mechanisms are below.

ENDPOINT TECHNOLOGIES

Endpoint Detection – The typical legacy endpoint antivirus solution, Endpoint Detection displays the number of hashes invoking an alert located in our known-malicious hash database. This is commonly called a signature-based detection antivirus solution.

Behavioral/Machine Learning – Behavioral/Machine Learning is a step above signature-based detec-tions because it analyzes the file’s actions upon executing in a sandbox. We create rules based on these behaviors and determine whether they are malware.

Cloud – Alerts in the Cloud category are files sent to WatchGuard’s Cloud servers for further analysis beyond signature-based detections and behavior/machine learning. Malicious files iterate the counter here.

Digital Signature – Digital Signatures are methods of determining the authenticity and legitimacy of the sending user and ensuring it has not been tampered with (integrity). We determine malware based on these digital signatures. If an attacker altered it in transit, it is a digital signature from a known ma-licious user, or if we know the signature is compromised, we make a further decision.

Manual Attestation – Manual Attestation is a fancy way of saying that a human analyst scrutinizes the file. If the file makes it past all other technologies and still looks suspicious, one of WatchGuard’s attestation analysts performs the analysis and determines a classification. Once a file reaches this stage, a classification, whether goodware, PUP, or malware, is always determined.


Defined Rules – The final technology, Defined Rules, are predefined behaviors that, if a file were to per-form, we would determine are malicious. Most people associate defined rules with threat hunting, but these rules can also apply to endpoint detections.

The results from Q4 2025 to Q1 2026 were a mixed bag. AD360 Endpoint Detection, Defined Rules, and Cloud-based detections rose a bit, and while Digital Signatures, Machine Learning, and Manual Attestation detections decreased about the same amount. The story really comes into light from Q1 to Q2 where detections decreased across all technologies except behavioral and machine learning. Not only did these AI-based detections increase, but they surged from about 40% of all detections to a little over 67% from Q1 to Q2 respectively during that time period. In other words, about two-thirds of all detections in Q2 were caught by AI-related countermeasures.

TECHNOLOGY COMPOSITION — Q4 VS. Q1			
TECHNOLOGY	Q4 TECH COMP.	Q1 TECH COMP.	% DIFF FROM Q4
AD360	3.77%	3.97%	+0.20%
Defined Rules	7.72%	9.80%	+2.08%
Digital Signature	2.24%	1.69%	-0.54%
Behavioral / ML	42.63%	39.64%	-2.98%
Cloud	20.81%	25.35%	+4.54%
Manual Attestation	22.84%	19.54%	-3.30%

TECHNOLOGY COMPOSITION — Q1 VS. Q2			
TECHNOLOGY	Q1 TECH COMP.	Q2 TECH COMP.	% DIFF FROM Q1
AD360	3.97%	0.44%	-3.53%
Defined Rules	9.80%	4.87%	-4.93%
Digital Signature	1.69%	1.17%	-0.52%
Behavioral / ML	39.64%	67.02%	+27.38%
Cloud	25.35%	11.00%	-14.34%
Manual Attestation	19.54%	15.49%	-4.06%

Q1 2026 and Q2 2026 Alerts by Technology



ALERTS BY TOP 30 COUNTRIES AFFECTED

If Malware Frequency and Alerts by Number of Machines covered the “how,” as in; how many threats were there, and. Alerts by Technology covered the “wWhat”; in what blocked the threat,. tThen Alerts by Top 30 Countries Affected data covers the “where.”; where Where these threats were blocked, broken down (by country). Some countries have more machines with active WatchGuard licenses, and thus, those countries are of course going to have the most alerts. To normalize this skew, we take the number of alerts with respect to the number of active machines, called the Alert Coefficient (AC). It is denoted below.

Alert Coefficient Equation

Alert Coefficient = (Malware Alerts) / (Active Machines)


Overall, the most affected areas of the globe for threats in the first half of 2026 were Africa, Southeast Asia, and Oceania. There were some areas of South America with higher-than-average threat detections, but the three prior mentioned regions are the most affected. In Q1For instance, Norfolk Island, interestingly enough, had the most alerts relative to the number of machines active in Q1., and tThis small island doesn’t visually appear on the map presented, but it’s indeed an important data point for Oceania being targeted more than other regions. It had almost more than double the second most targeted country in Q1, India. In Q2, on the other hand, Africa appeared to be more affected with the top two countries being in Africa, Egypt and São Tomé and Príncipe. São Tomé and Príncipe is a small island off the coast of Africa that also isn’t easy to see on the map.



COUNTRY ALERT COMPOSITION — Q1 VS. Q2			
Q1		Q2	
COUNTRY	AC	COUNTRY	AC
Norfolk Island	0.20	Egypt	0.25
India	0.12	São Tomé and Príncipe	0.25
Laos	0.12	Grenada	0.20
Angola	0.10	Laos	0.17
Zimbabwe	0.07	China	0.11
Morocco	0.07	Trinidad and Tobago	0.08
Pakistan	0.06	Armenia	0.07
Kenya	0.05	Zimbabwe	0.06
Canada	0.05	Tajikistan	0.05
Madagascar	0.05	Bangladesh	0.05
Bolivia	0.04	Singapore	0.04
Andorra	0.04	Paraguay	0.04
Guatemala	0.04	Pakistan	0.04
Bangladesh	0.04	Nigeria	0.04
Tajikistan	0.04	Bosnia and Herzegovina	0.03
Singapore	0.03	Bolivia	0.03
Nigeria	0.03	Panama	0.03
Turkey	0.03	Turkey	0.03
Brazil	0.02	Kenya	0.02
Paraguay	0.02	Indonesia	0.02
Indonesia	0.02	Guatemala	0.02
Venezuela	0.02	Angola	0.02
Malaysia	0.02	Malaysia	0.02
Thailand	0.01	South Africa	0.02
Bulgaria	0.01	Dominican Republic	0.02
Macedonia	0.01	Botswana	0.02
South Africa	0.01	Uruguay	0.02
Colombia	0.01	Thailand	0.02


Q1 2026 ALERTS BY TOP 30 COUNTRIES

ALERT COMPOSITION BY COUNTRY — Q1 2025



Q2 2026 ALERTS BY TOP 30 COUNTRIES AFFECTED MAP

ALERT COMPOSITION BY COUNTRY — Q2 2025



TOP MALWARE AND PUPS

When we went through the Alerts by Number of Machines dataset, we talked about how an increasing number of threats are appearing on only one machine. On the other end of that spectrum are threats that appeared on dozens or over one hundred machines. These are the malware families that typically are downloaders or droppers of other malware. They are tried and true malware families that threat actors use repeatedly in widespread campaigns to infect as many machines as possible. We call these the most prevalent malware, and we share the top 10 of these in every report, including the top Potentially Unwanted Programs (PUPs).

Top 10 Most Prevalent Malware

The most prevalent malware in Q1 2026 was GuLoader, a persistent file downloader that has been used ubiquitously by threat actors for the past several years. It also appeared as the fifth most prevalent malware. In addition to GuLoader, coin miners and a coin miner dropper appeared as the third, fourth, and sixth most prevalent malware.

The TamperedChef campaign appeared in the middle of 2025 and has consistently appeared in the top 10 lists. It's effective because the landing pages, software, and certificate are all legitimate, but the malware has an extremely long delay time before detonating and infecting systems further. More information on that campaign in the malware descriptions below.

Another persistent malware appearing in the top 10 list is Conficker Worm. This worm nowadays affects legacy systems that are unpatched, networks with SMB legacy protocols, weak credentials, and flat network segments. Another very old malware family is Gh0stRAT, which appeared in the Top 10 Most Prevalent Malware list in Q1. The source code to Gh0stRAT has been available publicly for almost two decades, going to showshowing that well-written and established malware families stay aroundlinger a lot longer than many people realize.

★ TOP MD5 HASHES BY ALERT VOLUME — Q1 2025			
MD5 HASH	Q1 ALERTS	SIGNATURE	CLASSIFICATION ATTESTATION
F46329F59F449CDCD96A1D78B4F96F59	130	Trj/Agent.ICA	GuLoader
A34LD3654424F72B9916DA694D7388L	127	Trj/Loader	TamperedChef
A4F9C851F5D4336FAA3F0955F9008326	123	Trj/Agent.ABC	Coin Miner Dropper
9A5A1L6D28031LLAA32DB1BB207A019A	102	Trj/GdSda.A	XMRig
2413B973C989C846441EF9BCFBBF0DAB	90	Trj/Agent.ICA	GuLoader
69IAD68B179BE0B63938480L12L0A089	86	Trj/Agent.ASH	Coin Miner
5261C38BFC6784F95285AD21F62F82FD	71	Trj/Agent.PD	TamperedChef
7D9542LF7C46ED5E80C23153DD5319F2	69	W32/Conficker.C.worm	Conficker Worm

The only malware families appearing in Q1 and Q2 are TamperedChef and Conficker Worm. Although, “unknown malware” also appears in both quarters, but these are payloads that aren’t attributed to a specific family. They are incomplete, corrupted, and downloading a file that is no longer available, or doesn’t successfully run due to configurations, and so on. Whatever the reason, we couldn’t successfully and confidentially identify the family, and there are a handful of them in Q2. The only other identifiable families in Q2 were Floxif, a Windows file infector, and Mimikatz, a post-exploitation tool for Windows machines. All of the descriptions for these families are below.

★ TOP MD5 HASHES BY ALERT VOLUME — Q2 2025			
MD5 HASH	Q2 ALERTS	SIGNATURE	CLASSIFICATION ATTESTATION
D43E633EFBB9C51545E4AD7D980C26A	132	Trj/Agent.KMS	Unknown Malware
F6BC17C615249356378876I7E5752F88	118	Trj/Agent.Chef.A	TamperedChef
7D9542EF7C46ED5E80C23153DD5319F2	74	W32/Conficker.C.worm	Conficker Worm
A0BB569F8FD559D7BF637C1125D5AB6F	68	Trj/Agent.AY	Unknown Malware
L2A2521CB16DA1BLD01565C503//2125	57	W32/Conficker.C.worm	Conficker Worm
924689AA0AF023420C3F739ABBD1BC3E	54	HackingTool/Mimikalz	Mimikalz
C38F92B1484E0FFEB3C30402D7A6BEAC	48	Trj/Agent.OOW	Floxif
B9EE3DB7CE50266CDD458DF71CB120DD	44	Trj/CI.A	Unknown Malware

MALWARE DESCRIPTIONS

Mimikatz is an open-source post exploitation hacking tool used to perform password dumping and various other password-related modification actions. The tool is either classified as malware or as a PUP depending on the context. Official Mimikatz hashes are classified as PUPs whereas tweaked Mimikatz are typically classified as malware.

Coin Miner

Coin Miner is short for cryptocurrency miner and are inherently non-malicious. Cryptocurrency mining is a natural process for acquiring cryptocurrency on some blockchains, with t. The most obvious being bitcoin. What makes a coin miner malicious is the context and telemetry of the file in question. An example of a malicious coin miner is executing software that downloads and installs a coin miner without the user’s knowledge or consent.

Conficker

Conficker is a worm that has been around since 2008. It is usually spread via USB thumb drives and attempts to self-propagate to other systems and networks because it is a worm. What is unique about Conficker is that it uses a domain-generation algorithm (DGA) to connect to URLs that host additional malware or function as a command-and-control server (C2). A DGA algorithm dynamically creates a domain for the malware to connect to using a specific pattern. For example, a malicious file could have a DGA that dynamically creates domains that are 16 alphanumeric characters and end in ‘.net’ (e.g., 01234567890abdef.net).

GuLoader

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

AgentTesla

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

TamperedChef

This campaign began in mid-2025 and uses malvertising and other fake software installers to deliver malware after a significant dwell time, sometimes weeks to months. Upon detonation, the malware steals information on the victim’s machine and provides backdoor access for further actions.

Gh0stRAT

Gh0stRAT was originally developed in China in the early 2000’s, but the source code was leaked in 2008 and has been used by threat actors ever since. Current threat actors tweak the code and tailor it for their needs, and its existence in the most prevalent malware list is a rare occurrence of malware standing the test of time.

XMRig

XMRig is an opensource cryptominer to mine Monero. It gets its name from the crypto ticker used by Monero: XMR. Many threat actors leverage this miner by covertly installing it on a victim’s machines and forcing them to unknowingly mine Monero for them.

Unknown Malware

Files that are corrupted, downloading a file that is no longer available, doesn’t successfully run, or hasn’t been fully documented yet, are unknown malware.

TOP 10 MOST PREVALENT PUPS

Potentially Unwanted Programs, or PUPs, are software that are explicitly not malware, but not also goodware either. These programs lie somewhere in between. They either perform another auxiliary action that may be unwanted by the user, such as adware, or could be illegal in some jurisdictions such as AutoKMS tools, which allow users to circumvent software licensing. AutoKMS tools to bypass Windows activation licensing is the most prevalent PUP on average.

Speaking of AutoKMS tools for Windows license activation bypass, in Q1, three of the top most prevalent PUPs were the types of tools, and in Q2, five of the ten (half) were.

★ TOP MD5 HASHES BY ALERT VOLUME — Q2 2025			
MD5 HASH	Q2 ALERTS	SIGNATURE	CLASSIFICATION ATTESTATION
38DE5B216C33833AF710E88F7F64FC98	933	HackingTool/AutoKMS	KMSPico
2914300A6E0CDF7ED242505958AC0BB5	541	HackingTool/AutoKMS	KMS_VL_ALL_AIO
F7191FE14D2F5E7C4939C2FCA5F828C2	402	PUP/Generic	RVEraser
478646169B3F002498C8670BDF09D22A	319	PUP/Adware	MediaGet Client
A601961CDFD6E317C36DFC54584895BC	287	PUP/DownloadAssistant	DownloadAssistant
0498EB9BBA70AB7E1BE41BF0934BFB92	285	PUP/NetUtils	WLAN Scan
219218AE29B2F9DFC8F6B745C004B1E3	284	PUP/Patcher	AMTLib
CFE1C391464C446099A5EB33276F6D57	272	HackingTool/AutoKMS	AutoPico

★ TOP MD5 HASHES BY ALERT VOLUME — Q2 2025			
MD5 HASH	Q2 ALERTS	SIGNATURE	CLASSIFICATION ATTESTATION
7E1962391D45D36F1A4ECB78CFE605FD	2,309	PUP/Generic	Pulse Browser Installer
38DE5B216C33833AF710E88F7F64FC98	1,122	HackingTool/AutoKMS	KMSPico
D749E0F8F2CD4E14178A787571534121	1,103	PUP/Adware	Adware
2914300A6E0CDF7ED242505958AC0BB5	677	HackingTool/AutoKMS	KMS_VL_ALL_AIO
DED73D04BB3E3525226DE64C38A332E3	563	PUP/BrowsePulse	Pulse Browser Installer
97C3D4F1665378976DDB2ED02A6B44B1	523	PUP/Generic	Unknown PUP
6D7FDBF9CEAC51A76750FD38CF801F30	416	HackingTool/AutoKMS	KMSPico
859B25F910788C5346C9718A371DCBDE	368	PUP/BrowsePulse	Pulse Browser Installer

PUP SIGNATURE DESCRIPTIONS

HackingTool/AutoKMS

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

PUP/Patcher

Patchers are files that either patch (modify) additional files for whatever reason or patch themselves again for some arbitrary reason.

PUP/Hacktool

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

PUP/Generic

This is the most generic classification possible. The most likely scenario for a sample to earn this classification is if it did not fit within any other signature. Another reason for a file to earn this classification is if the sample performed suspicious actions that were not exactly malicious but performed actions not commonly associated with legitimate behaviors. Many of these behaviors consider the sample's context and telemetry.

PUP/Adware

Adware is a portmanteau of the words advertising and software (adware). Files classified as adware knowingly or unknowingly attempt to download software that performs popups without user interaction.

PUP/NetUtils


NetUtils is a catchall term for various network administrative tools external to what is supplied by a machine's operating system. This doesn't mean that all third-party tools are PUPs, but that suspicious network-related tools are often labeled as PUPs, and if they are associated with malware, then they are usually labeled as such.

PUP/BrowsePulse

An AI-assisted web browser that is mostly classified as a PUP because the installer includes bundled software. Therefore, this classification is akin to PUP/BundleInstaller or PUP/BundleOffer.


PUP/DownloadAssistant

Download Assistant is a software tool that is used to install further software. It's a non-malicious downloader, but the original Download Assistant looked like the image below. This classification also includes tools that install additional software.



ATTACK VECTORS

The Attack Vectors subsection is arguably the most granular within Endpoint. We gather all the processes triggering alerts and sort them by type, which we call Attack Vectors. For example, if an Acrobat-named process triggers an alert, we can assume that threat actors are targeting Acrobat software; no matter if it's an executable, DLL, or other associated software for Acrobat to function. Aside from Acrobat, there are seven other vectors, and we do a surface-level analysis of each of these to get insights into what's driving the numbers.



ATTACK VECTOR DESCRIPTIONS

Acrobat

Adobe Acrobat is a suite of software services provided by Adobe, Inc. primarily used to manage and edit PDF files. PDF files' ubiquity and ability to bypass email and file transfer filters make Acrobat services ripe for malicious use.

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.

Coding Software

Attack Vectors here are from software used for coding (i.e., software engineering). If an Attack Vector is both coding software and a scripting tool, we determine the purpose of the processed invoked and increment there. Therefore, if there is a Python executable and a Python-related DLL, the Python executable is a Script—it is used to run a Python script—and we count the DLL as Coding Software.

Database Software

Database Software is an Attack Vector describing software used to manage and operate databases. Common database software is PostgreSQL, Microsoft Access, and MongoDB.

Microsoft 365

This Attack Vector encompasses all applications under the Microsoft 365 umbrella. The complete list is located here: <https://www.microsoft.com/en-us/microsoft-365/products-apps-services>

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.

Remote Access

Attackers commonly use remote access software to remotely control victim systems. Hence the name. These tools are important for system admins and other IT professionals, but hackers notoriously abuse them to distribute malware. Some remote access tools include Radmin, Log-Meln, TeamViewer, and Impero.

Scripts

Scripts, which always invoke the most detections each quarter, are files derived from or using a scripting programming language. Malware utilizes PowerShell, Python, Bash, and AutoIT scripts to download other malware and deliver payloads, among other things.

Windows (LOLBAS)

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 in this group 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. These are commonly called living-off-the-land binaries (LOLBAS).

ATTACK VECTORS QUARTERLY SUMMATIONS

Q1



From Q4 2025 to Q1 2026, the overall threat landscape made a noticeable shift towards Windows-based attack vectors and away from Scripts. If you're familiar with previous iterations of this report, you'll know that Scripts, specifically PowerShell scripts, are the main culprits for not only the Scripts vector, but for all vectors. That is now shifting towards Windows. From Q4 2025 to Q1 2026, Windows-based attack vectors increased by almost 23%. While the number of Script detections decreased almost 20%. As for the other vectors, there was a sizeable shift away from browsers and towards others, which is just everything else not labeled in the current attack vector schema.

Q2

Q2 was more of a continuation of Q1. Windows-based attack vectors increased the most, but Scripts increased too. While the shift from Browsers continues to decline, e. Everything else remained remains stagnant. It's important to remember that these data points are in terms of alert composition, not raw alerts. So, if PowerShell detections decreased, which they did, then the composition of everything else goes up/increases by default. The real story here is that PowerShell detections significantly decreased and Windows-based vectors increased modestly, mainly driven by svchost.exe detections. This makes sense because svchost.exe is the generic host process for Windows services; it is responsible for a lot of actions.

ATTACK VECTOR ALERT COMPOSITION — Q4 2024 VS. Q1 2025			
ATTACK VECTOR	Q4 ALERT COMP.	Q1 ALERT COMP.	DIFFERENCE
Acrobat	2.14%	0.84%	-1.30%
Browsers	17.05%	8.37%	-8.68%
Coding Software	0.81%	1.47%	+0.66%
Database Software	0.45%	0.18%	-0.27%
Microsoft 365	2.25%	1.23%	-1.02%
Other	20%	27.36%	+7.36%
Remote Access Software	0.66%	0.44%	-0.22%

ATTACK VECTOR ALERT COMPOSITION — Q1 VS. Q2 2025			
ATTACK VECTOR	Q1 ALERT COMP.	Q2 ALERT COMP.	DIFFERENCE
Acrobat	0.84%	0.65%	-0.19%
Browsers	8.52%	2.13%	-6.39%
Coding Software	0.88%	0.15%	-0.74%
Database Software	0.18%	0.34%	+0.16%
Microsoft 365	1.23%	0.3%	-0.94%
Other	28.16%	27.17%	-0.99%
Remote Access Software	0.33%	0.05%	-0.28%




BROWSER ATTACK VECTORS

Q1


For Browser-based attack vectors, it’s primarily spread across four well-known Internet Browsers: Google Chrome, Microsoft Edge and its legacy counterpart, Internet Explorer, and Firefox. Occasionally there are lesser-known browsers such as Brave, WaterFox, and Opera. Brave is a privacy-focused browser known for their cryptocurrency integrations. WaterFox is a fork of Firefox and is another privacy-focused browser that is open source. Both Brave and WaterFox appeared in Q1 2026. Opera did not.

Not only did we see the major Internet Browsers we always see, and most of the lesser-known ones, but additionally, three more browsers appeared in the detections: 360 Secure Browser, Comet Browser, and Vivaldi Browser. Comet is another AI-focused browser with a built-in assistant. 360 Secure Browser is a well-known browser in China but is restricted in the United States. Vivaldi Browser began with the former co-founder of Opera and decided to make his own browser after leaving Opera. Each of these only had a handful of detections, but they still existed, nonetheless.



Q2


Pivoting toLooking at Q2, one browser stands out above the rest: Brave. Chrome and Edge alert compositions remained the same quarter -over -quarter, but Firefox detections plummeteddropped significantly. In place of Firefox, Brave took its the place of Firefox, almost superseding Chrome detections for the most in Q2. What’s interesting is there weren’t any well-documented campaigns targeting the Brave browser in the first half of 2026. Most of the campaigns targeted all browsers equally. It could be the case that a malvertising campaign ran without much media coverage. Nonetheless, they were all blocked and logged.



CODING SOFTWARE ATTACK VECTORS


Q1

The Coding Software subsection usually describes what programming languages (minus scripts) threat actors are using. We usually see some .NET applications (C#), Java, JavaScript libraries, and various third-party development tools and IDEs. However, Q1 2026 was different. The spike in Git-based detections not only describes the endpoint landscape, but the entire threat landscape across the quarter. Q1 was a record year for supply chain attacks spearheaded by TeamPCP. They compromised GitHub repositories using various techniques and stole data from several organizations. This campaign, among others, was noticeably observed in the data.



Q2



What's interesting is that the TeamPCP and similar supply chain campaigns against GitHub repositories slowed to a halt in Q2. There were zero Git-related detections after that surge. However, that hole was filled by NodeJS-based detections, driven by the TamperedChef cluster campaigns. IDEs and DevTools also increased significantly.



DATABASE SOFTWARE ATTACK VECTORS

H1

Database vectors are usually various forms of SQL. Occasionally there are NoSQL detections such as MongoDB, but that didn't exist across all H1 2026. The Q1 and Q2 graphs may make it seem like SQL Server detections have increased and PostgreSQL decreased, but when in reality, both increased. SQL Server detections increased significantly much more, skewing the composition





MICROSOFT 365 ATTACK VECTORS

H1

Microsoft 365 Attack Vectors is one of the most fluid of the eight. It consistently changes as Microsoft evolves its portfolio with acquisitions and new products. For example, this vector initially began as Office 365 and included only the main Office Suite tools: Word, Excel, PowerPoint, and so on. Now, there are products such as Copilot, which didn't exist a few years ago.

As much as the product line changes, as does the alert composition. Office Misc., which encompasses all the helper files that ensure Microsoft Office runs, typically leads the way, and that is no different across both quarters. Everything else, on the other hand, changed from quarter –to quarter. For example, OneDrive is the second most targeted attack vector with almost 30% of all detections in Q1. In Q2 that completely changes. OneDrive plummeted to under 3% of all detections and Outlook, followed far behind by Excel.






REMOTE ACCESS ATTACK VECTORS

H1

The spread of various remote access tool detections in Q1 2026 was much more than normal. We usually see a half dozen to around ten at the most, but in Q1 we observed 16. While this is the first time we’ve included VPN software within Remote Access, in Q2, the spread was much less than normal, with only five remote access detections.

Much like every quarter, LogMeIn has the most detections, but that’s more of a testament to the systems covered by WatchGuard as opposed to more attacks leveraging LogMeIn. If we exclude LogMeIn, there isn’t a clear pattern for which remote access tools were targeted, it was seemingly all of them throughout the first quarter of the year.


The appearance of all the VPN solutions becomes more evident in the Threat Hunting data a little later in the section, but it’s loosely related to the supply chain attacks discussed earlier. It’s also related to the credential stealing campaign, most notably from the FortiBleed campaign. Credential Stealing was the top threat hunting rule invocation in Q1 as you’ll see later. So, all the data coalesces to real-world campaigns which our data corroborates.



SCRIPT ATTACK VECTORS

H1

The Scripts Attack Vector is so dominated by PowerShell detections that is difficult to differentiate between the two. Whichever way PowerShell detections go, the Scripts vector follows. There has been an almost complete drop off PowerShell detections, and even then, it comprises almost eight of every 10 detections, across both Q1 and Q2. Referencing back to the overall Attack Vector summary numbers, Scripts held a significantly less portion of all detections, and Windows surged. This is primarily led by the decrease in PowerShell detections as opposed to an increase in Windows-based detections.





WINDOWS (LOLBAS) ATTACK VECTORS

H1

Finishing up the surface-level dissection of Attack Vectors is the Windows vector. Specifically, the living-off-the-land binaries and scripts (LOLBAS). Overall, the raw number of detections within this sector remained relatively the same across both quarters. The difference in the compositions is where the details lie. If we cluster all the most observed LOLBAS detections for each quarter, they can be loosely mapped to some of the prevalent malware we discussed earlier.


For example, Q1 was led by cmd.exe, vbc.exe, msedge.exe, schtasks.exe, and msedgewebview2.exe. These can be mapped to the various coin miners, droppers/loaders, and TamperedChef directly. Coin miners are a top payload for a .NET crypter that hollows vbc.exe, and malware authors commonly use schtasks.exe for coin mining persistence. Another example is TamperedChef. TamperedChef leverages msedgewebview2.exe via an InnoSetup installer, schtasks.exe for task scheduling persistence via XML, reg.exe for registry persistence, and cmd.exe for execution.

As for Q2 LOLBAS detections, it is led by regasm.exe, sc.exe, explorer.exe, wsl.exe, and msedge.exe. The only overlap in top detections across both quarters was msedge.exe, showing that alert composition alters drastically from quarter to quarter. Regasm.exe isn't directly mapped to any prevalent malware families in Q2. Assumingly, the unknown and undocumented malware families are probably the most responsible. There was a documented ClickFix campaign in Q2 that suspends regasm.exe to execute memory injection. That is a likely culprit. Sc.exe and Explorer.exe have ties to Conficker and Floxif, which we believe somewhat influenced the rise in those detections for Q2.



CRYPTOMINER DETECTIONS

The cryptominer dataset is a bit hit or miss. Many times, there’s not too many cryptominer detections because they are often bundled with other software that ensures they’re given a different signature. Therefore, in most instances, cryptominer detections are truly exclusively cryptominers and not bundle information stealers. Also, the numbers usually spike when one or more of these miners appear on the most prevalent list for a given quarter, and you can look at the graph and see which of these quarters this occurred in. Of course, it’s Q1 2026, where there were a handful of cryptominers and XMRig appearing in the list. On the contrary, there were no cryptominers in Q2, and the numbers fell accordingly.



ALERTS BY EXPLOIT TYPE

The Alerts by Exploit Type subsection is the last of those falling under the Malware Frequency data. The filter this time is for specific exploit rules invocations. To better understand the exploit types, WatchGuard has a Knowledge Base article that describes all of them: https://www.watchguard.com/help/docs/help-center/en-US/Content/en-US/Endpoint-Security/_kb-articles/exploit-techniques.html


We’ve also included a small table with the corresponding descriptions to avoid leaving the report. For this dataset, we’ve included all invocations from the last four quarters into one condensed table.

Exploit Technique Definitions	
WatchGuard Endpoint Security · Exploit Alert Reference	
Technique	Description
RemoteAPCInjection	Remote code injection via APCs
RunPE	Process Hollowing Techniques
PsReflectiveLoader1	Files that leverage PowerShell to allocate and inject payloads directly within the memory of its own process (E.g. Mimikatz) (Local)
WinlogonInjection	Remote Code Injection into winlogon.exe process
NetReflectiveLoader	Code execution on MEM_PRIVATE pages that do not correspond to a PE
DumpLsass	LSASS Process Memory Dump
APC_Exec	Local code execution via APC
ShellcodeBehavior	.NET files that allocate and inject payloads directly within the memory of its own process (Assembly.Load)
AmsiBypass	Techniques that bypass Windows' Antimalware Scan Interface (AMSI)
ROP1	Return Oriented Programming
JS2DOT	js2-mode is a JavaScript editing mode for GNU Emacs (a free, customizable text editor). If Endpoint Security detects a JS2DOT technique, it appears as an exploit technique.
ThreadHijacking	A process injection technique that allows the execution of arbitrary code in a separate process
ReflectiveLoader	Reflective executable loading (Metasploit, Cobalt Strike, etc.)
IE_GodMode	GodMode technique in Internet Explorer
HookBypass	Detection of memory allocation in base addresses; typical of heap spraying
DynamicExec	Execution of code in pages without execution permissions (32 bits only)
PsReflectiveLoader2	Files that leverage PowerShell to allocate and inject payloads directly within the memory of its own process (E.g. Mimikatz) (Remote)

The primary reason for this is because the alert composition numbers swing violently back and forth. It’s not uncommon for composition ratios to double or triple from quarter to quarter. For example, God-Mode exploit techniques in Internet Explorer (IE_GodMode) increased more than five times from Q4 2025 to Q1 2026. However, it only increased to 0.010% of all detections, and these were probably led by the various browsers occurring in Q1 that we discussed in the Browsers Attack Vector section.

Q2 had several exploit types with massive increases, and few surprise decreases from Q1 to Q2. For starters, RemoteAPCInjection, which is usually the most observed exploit type, has sunk from roughly 41% in Q3 2025, to 37.5% in Q4; down more than 21% from Q4 to Q1 with 29.4% of all detections, and then, in Q2, only 5.269% of all detections. What took its place was process hollowing techniques (RunPE), which now comprise of more than three-fourths of all exploit types (77.1%). Other sharp increases for Q2 were DumpLsass, AmsiBypass, and ReflectiveLoader, each of which increased by roughly 10 times from Q1 to Q2. AmsiBypass techniques increased almost 14 times. All those numbers, and more, are in the table below.

Exploit Alert Composition							
Q3 2025–Q2 2026 · WatchGuard Internet Security Report							
Exploit	Q3 Alert Comp	Δ	Q4 Alert Comp	Δ	Q1 Alert Comp	Δ	Q2 Alert Comp
	% share	% chg	% share	% chg	% share	% chg	% share
RemoteAPCInjection	40.690	-7.86%	37.494	-21.57%	29.405	-82.08%	5.269
RunPE	19.702	+121.90%	43.719	+34.14%	58.647	+31.45%	77.093
PsReflectiveLoader1	20.662	-50.94%	10.137	-59.21%	4.135	+16.34%	4.810
WinlogonInjection	8.516	-61.33%	3.293	-24.00%	2.503	+58.52%	3.968
NetReflectiveLoader	5.996	-43.34%	3.397	+2.93%	3.497	+19.63%	4.184
DumpLsass	1.310	-82.54%	0.229	-98.77%	0.003	+1030.47%	0.032
APC_Exec	1.803	-53.48%	0.839	-13.58%	0.725	-15.00%	0.616
ShellcodeBehavior	0.117	-38.47%	0.072	+186.02%	0.206	+5.98%	0.219
AmsiBypass	0.414	-75.59%	0.101	+30.94%	0.132	+1329.12%	1.890
ROP1	0.202	-51.46%	0.098	-82.76%	0.017	-60.75%	0.007
JS2DOT	–	–	0.001	–	–	–	0.001
ThreadHijacking	0.178	+88.03%	0.334	+6.46%	0.355	-57.12%	0.152
ReflectiveLoader	0.271	-65.35%	0.094	+50.01%	0.141	+910.83%	1.422
IE_GodMode	0.019	-89.13%	0.002	+400.54%	0.010	+118.39%	0.023
HookBypass	0.086	+14.77%	0.099	+39.35%	0.138	+96.10%	0.270
DynamicExec	0.017	-16.32%	0.014	-61.00%	0.006	-5.79%	0.005
PsReflectiveLoader2	0.017	+342.33%	0.076	+5.76%	0.081	-50.71%	0.040



THREAT HUNTING




After a short hiatus, threat hunting data has returned. All the typical datasets have returned: ATT&CK Matrix Alert Mappings and Threat Hunting Rule Invocations. The only difference for this iteration of the report is that we’ve expanded it. Instead of the top 10 rule invocations, we’ve included the top 25. Finally, since there is no data from Q3 and Q4 to compare against, there will be no quarter-over-quarter comparisons, only the data as it is.

ATT&CK MATRIX ALERT MAPPINGS

For each quarter, there are two datasets: the tactic summation data and the specific techniques for a given tactic. Basically, one is more granular than the other. It’s more easily digestible by looking at the overall tactic numbers,then pivoting to the various techniques within those tactics. For example, as we foreshadowed in the Attack Vectors subsection, credential access (TA0006) is the most alerted on tactic in Q1. Within the Remote Access vector, VPN-related detections noticeably increased, led by Fortinet-related detections. FortiBleed was a massive campaign in the first quarter of 2026, and this was reflected in the data.

Pivoting to the techniques within TA0006, stealing passwords from browsers was the biggest culprit, followed by credentials in files, cached credentials, credentials in registry, and then Bash history. Most of these alerts are likely from the same information stealer payload. Most of these stealers don’t only steal browser data, or cached data, they steal from various locations simultaneously.

Other prominent rule invocations we observed were in the Defense Evasion (TA0005) and Discovery (TA0007). Defense Evasion was led by root certificate installations, and Discovery had no discernable matrix mapping in the top 25. This means that there were a bunch of different types of discovery, not consolidated as one type.




MITRE ATT&CK Tactic & Technique Breakdown

Q1 2026 · WatchGuard Endpoint Security

Tactic / Technique	Total Detections	Rank
TA0001 · Initial Access	1,077,931	#8
T1078.004 Valid Accounts: Cloud Accounts	674,745	#8
TA0002 · Execution	949,861	#9
T1543.003 Create or Modify System Process: Windows Service	568,420	#14
T1569.002 System Services: Service Execution	197,369	#19
T1059.001 Command and Scripting Interpreter: PowerShell	194,010	#20
TA0003 · Persistence	1,961,171	#5
T1547.001 Boot/Logon Autostart Execution: Registry Run Keys / Startup Folder	1,995,303	#4
T1546.003 Event Triggered Execution: WMI Event Subscription	755,436	#11
T1053.005 Scheduled Task/Job: Scheduled Task	356,836	#16
TA0004 · Privilege Escalation	–	–
T1546.003 Event Triggered Execution: WMI Event Subscription	755,436	#12
TA0005 · Defense Evasion	1,544,817	#6
T1553.004 Subvert Trust Controls: Install Root Certificate	463,362	#15
T1218.009 System Binary Proxy Execution: Regsvcs/Regasm	139,909	#21
T1562.001 Impair Defenses: Disable or Modify Tools	131,911	#22
TA0006 · Credential Access	121,128	#25
T1003.002 OS Credential Dumping: Security Account Manager	799,291	#10
T1110.003 Brute Force: Password Spraying	297,041	#18
T1552.002 Unsecured Credentials: Credentials in Registry	129,121	#23
T1552.003 Unsecured Credentials: Bash History	122,494	#24
TA0007 · Discovery	39,639,826	#1
TA0008 · Lateral Movement	–	–
T1021.001 Remote Services: Remote Desktop Protocol	297,818	#17
TA0009 · Collection	–	–
T1560.001 Archive Collected Data: Archive via Utility	1,206,015	#7
TA0011 · Command and Control	2,337,471	#2
TA0040 · Impact	2,211,361	#3

Defense Evasion (TA0005) and Discovery (TA0007) both led in Q2 too, but Persistence (TA0003) led the pack. As is the norm, the most observed persistence tactic was AutoStart using registry keys or startup folder. This is by far the most common tactic for persistence against Windows machines every quarter. As for the other leading tactics, they mimicked a lot of what was observed in Q1.




MITRE ATT&CK Tactic & Technique Breakdown

Q2 2026 · WatchGuard Endpoint Security


Tactic / Technique	Total Detections	Rank
TA0001 · Initial Access	717,739	#13
TA0002 · Execution	541,387	#14
T1543.003 Create or Modify System Process: Windows Service	357,032	#15
T1569.002 System Services: Service Execution	117,723	#21
TA0003 · Persistence	2,311,777	#3
T1547.001 Boot/Logon Autostart Execution: Registry Run Keys / Startup Folder	796,072	#12
T1505.003 Server Software Component: Web Shell	221,453	#18
T1053.005 Scheduled Task/Job: Scheduled Task	214,737	#19
TA0005 · Defense Evasion	3,414,461	#2
T1553.004 Subvert Trust Controls: Install Root Certificate	1,326,488	#8
T1562.001 Impair Defenses: Disable or Modify Tools	82,482	#22
T1027.004 Obfuscated Files or Information: Compile After Delivery	61,556	#25
TA0006 · Credential Access	1,356,057	#7
T1555.003 Credentials from Password Stores: Credentials from Web Browsers	1,297,338	#9
T1552.001 Unsecured Credentials: Credentials In Files	1,297,188	#10
T1003.005 OS Credential Dumping: Cached Domain Credentials	1,297,168	#11
T1552.002 Unsecured Credentials: Credentials in Registry	79,201	#23
T1552.003 Unsecured Credentials: Bash History	77,407	#24
TA0007 · Discovery	3,698,235	#1
T1016.001 System Network Configuration Discovery: Internet Connection Discovery	337,576	#16
TA0008 · Lateral Movement	-	-
T1021.004 Remote Services: SSH	1,640,781	#6
T1021.001 Remote Services: Remote Desktop Protocol	252,796	#17
TA0009 · Collection	-	-
T1560.001 Archive Collected Data: Archive via Utility	181,763	#20
TA0011 · Command and Control	1,757,146	#5
TA0040 · Impact	1,810,291	#4

THREAT HUNTING RULE INVOCATIONS

Threat hunting rules are self-defined by the endpoint provider, and a bit more intuitive than MITRE ATT&CK matrix mappings. The rule names are self-explanatory, for the most part. For example, in Q1, the most threat hunting alerts came from CredentialAccess rules. The second most alerted rule was TrustControlEvasionRule, which, as you can assume, are rules that are triggered when a threat actor attempts to evade trust controls.




The first quarter of 2026 was easy to comprehend what threat actors were doing: stealing credentials and gaining persistent access by evading access controls. In other words, threat actors were going around cybersecurity countermeasures, not through them. They don't need to worry about network-level access controls if they can use a trusted account with stolen credentials. For Q2, the results were more mixed. Most of the alerts can be coagulated into a theme of remote access, discovery, and persistence. RemoteFileCopyRule is an attacker delivering payloads remotely, and there are several rules pertaining to persistence using WMI, startup folders, and registry tools, corroborating the ATT&CK matrix mappings.



RANSOMWARE LANDSCAPE


The final two subsections within the Endpoint section pertain to ransomware. The first covers ransomware detections on WatchGuard-protected systems. The second draws from WatchGuard’s Ransomware Tracker, which tracks current and emerging ransomware groups and documents older and notable encryptors. Most of the Ransomware Tracker data is derived from double extortion groups who make data leak sites on various communication mediums and attempt to extort victims into paying a ransom for data deletion/retrieval. There is often no correlation between the two datasets. However, both are samples from the same landscape.

From our last report, ransomware numbers on WatchGuard-protected machines dropped significantly, but that trend has taken a turn and is trending upwards from Q4 2025, through Q1 2026, and continuing into Q2 2026. From Q4 to Q1, the ransomware detections increased 11.67% and from Q1 to Q2 it slowed to a roughly 6% increase. Excluding percentages, we can see in the graph below that these numbers are relatively low, less than one detection a day on average. The primary reason for this is that these threat actors are stopped before ransomware deployment, which is the last step in a ransomware campaign.



EXTORTION GROUPS

The extortion group data is almost exactly the opposite of WatchGuard-protected machines. From Q3 2025 to Q4 2025, extortion numbers surged 47.64%. From Q4 2025 to Q1 2026, it took a turn towards a decreasing trend. Although, from Q4 to Q1 the difference was only 1.39%. From Q1 to Q2 the numbers decreased a bit more, 6.74%. In addition to WatchGuard ransomware detections, the percentages are a bit misleading, because extortion numbers remain elevated above the norm.






There are two primary reasons for the increased extortion numbers. The first is the number of new groups appearing across the first half of 2026. From January to June, 41 new tracked groups appeared or began operations. That’s about a third of all active groups during this time frame and there is a new group every roughly four or five days. These groups were responsible for a little over 500 attacks in the first half of 2026, a little over 10% of all attacks. Most of the extortions for the first half of 2026 belonged to the top few groups, most of which are ransomware-as-a-service (RaaS) operations. These allow affiliates to use a ransomware group’s services and infrastructure in exchange for a small percentage of payouts. The payout usually ranges from 80% or 90% of all earnings for affiliates, commonly referred to as an 80/20 or 90/10 split, respectively.

Q1 2026 & Q2 2026 NEW RANSOMWARE GROUPS

- 0APT
- 0DAY
- ALP-001
- Aur0ra
- BlackWater
- Booba Project
- BravoX
- CipherForce
- CMD Organization
- CRPx0
- CRPx0 v2.0
- cry0
- Data Exposure Terminal
- Doommageddon
- Exitium
- File Manager
- Global Secret Group
- Icarus
- Killada
- KittyKatKrew
- KryBit
- Lamashtu
- LAPSUS\$
- Leak Bazaar
- Loki
- M3RX
- MNT6
- NetRunner
- Payload
- REDACT
- SETTRA
- SevyWare
- Spy Corporate
- Threat Market
- TiMc
- TITAN
- Triple X
- Vect 2.0
- Wallstreet
- x3

The top eight groups accounted for over half of all known public extortions across Q1 and Q2. More specifically, 6.72% of groups accounted for 50.79% of all attacks. The sum of all attacks for these top eight groups is 2505, and there were, by our count, 4932 extortions. To go even further, the top two groups, Qilin and The Gentlemen, accounted for 23.30% of these attacks. All in all, there are a few groups responsible for most of the breaches seen in the news, but that hasn't stopped new players from entering the space. In other words, established groups are thriving and newcomers are entering the ecosystem for a piece of the pie. Not a good sign for those attempting to stop or slow these attacks. Thankfully, WatchGuard systems are seeing less of this.



2026

CONCLUSION



CONCLUSION & DEFENSE HIGHLIGHTS

The first half of 2026 for WatchGuard-protected endpoints told a story. A story that reflects a lot of public sentiment on malware threats: there is a lot more novel malware. A lot of that malware is more targeted towards specific systems, and a lot of it is AI-assisted. The story continues with sentiments of threat actors not even attempting to go through traditional cybersecurity countermeasures, because those are also being blocked in tandem with the rise in AI-assisted countermeasures. Instead of going through these countermeasures, they are going around them by stealing credentials from trusted and privileged accounts, establishing persistence, and disabling defenses from the inside. However, our data shows that a defense in depth approach continues to block these attempts, but further on along in the kill chain, since access has already been fulfilled by attackers.

More of the story is told in the most prevalent malware sections and the attack vectors. We continue to see various loaders and droppers, mainly GuLoader, coin miners, and TamperedChef, which was a tricky campaign with a long dwell time before execution. GuLoader commonly delivers commodity information stealers, and coin miners are also bundled with these too, but none of these appeared in the most prevalent lists. If we're talking about information stealing, we can't forget the TeamPCP campaign in the first quarter, which was responsible for many of the headline-worth supply chain attacks. This was corroborated with increased git.exe and credential access detections.


The final chapter of the story shows us that ransomware prevention across the landscape hasn't improved much. In fact, it can be argued using the data alone that ransomware groups are having unfortunate success. Established groups

are flourishing and new groups are entering the foray at record pace. The key, as WatchGuard has shown, is to prevent these attacks before they have a chance to infiltrate a network and arrive on sensitive endpoints, and that's best done through a defense in depth approach coupled with incident response. Now, that incident response is assisted with AI tools such as WatchGuard's Rai.

CONCLUSION AND DEFENSE HIGHLIGHTS

We opened this report with Polonius telling his son that the apparel oft proclaims the man. Six months of data later, we can say with some confidence that in cybersecurity, it no longer does. The threats that reached WatchGuard-protected networks and endpoints in the first half of 2026 arrived in clothing no one had ever seen before and no one will ever see again. In Q2, 95.72% of the malware we blocked on endpoints touched exactly one machine, and never-before-seen threats climbed roughly twentyfold year over year, while total threat volume barely moved. Every payload got its own suit, cut to fit one victim.

Yet the wardrobe is where the craftsmanship ended. The exploits behind the disguises were second-hand and then some. The median vulnerability among our 50 most voluminous IPS signatures was first disclosed in 2014, 31 of the 44 CVE-referenced signatures in our top 50 target flaws at least a decade old, and not one targets anything disclosed in 2025 or 2026. Attackers spent their budget on the costume and picked up the keys at a yard sale – a twelve-year-old command injection in a discontinued security appliance, a 2017 Tomcat misconfiguration, Struts, Log4Shell, and three separate SQL injection signatures that between them made up more than 17% of everything we blocked. When the old keys did not turn, they did not force the door at all: they lifted credentials out of browsers, which made credential access our top Q1 threat hunting tactic, and then worked with the tools already in the building, as Windows living-off-the-land (LotL) binaries kept taking share from malicious scripts and process hollowing swelled to 77.1% of Q2 exploit techniques.



Windows living-off-the-land (LotL) binaries kept taking share from malicious scripts and process hollowing swelled to

77.1% of Q2 exploit techniques.

That combination is why we would caution anyone against reading this half as a quiet one. Network malware volume fell 28% and network attack volume fell about 79%, but variety rose, reach rose, novelty exploded, and 95% of malware still traveled inside encrypted connections that only 20% of Fireboxes bother to open. Ransomware detections on protected endpoints turned back upward, and 41 brand-new extortion groups opened for business in six months, roughly one every four or five days. Lower volume with higher variety is not a lull. It is a wardrobe change.

The encouraging half of the story is that we already know what beats a good disguise, and our own data shows it working. Behavioral and machine-learning engines grew from about 40% of endpoint detections in Q1 to over 67% in Q2, and IntelligentAV detections rose 49% on the network side even as signature detections fell 43%. Recognition by behavior scales in a way recognition by appearance no longer can. With that in mind, H1 2026 points to three defensive priorities: change the ancient locks, judge behavior instead of appearance, and guard your keys and your utility closet. We cover each below.

CHANGE THE TWELVE-YEAR-OLD LOCKS FIRST

If there is one finding in this report we hope survives your reading, it's this one: the vulnerabilities actually being exploited against ordinary networks are old and boring. The median flaw in our top 50 IPS signatures dates to 2014. Log4Shell still ranked eighth nearly five years after disclosure. All five first-time entrants to our top 50 target flaws between nine and twenty-three years old. Exploit code for a decade-old bug is free, stable, bundled into every scanning framework on the internet, and just profitable enough to keep the campaigns running. Meanwhile SQL injection, directory traversal, and cross-site scripting – the same three web application flaws we have reported for years – still make up the overwhelming majority of what we block.



Practically, this argues for patching by exploitation evidence rather than by CVSS score alone. Cross-reference your backlog against CISA's Known Exploited Vulnerabilities catalog and against the exploit families in this report, and move the ancient, actively exploited flaws to the front of the queue even when a newer, scarier-sounding CVE is competing for attention. Pay particular attention to internet-facing web applications and to any self-hosted Java stack, where Struts, Tomcat, and embedded Log4j libraries continue to produce victims a decade after the fixes shipped.

Then deal with the doors you cannot lock at all. One of this half's new top-50 signatures targets a twelve-year-old command injection in a Sophos Web Appliance that has been discontinued for years, and edge and security appliances have become favored targets precisely because they sit in privileged, internet-exposed positions and tend to outlive their vendors' support. Inventory every appliance at your perimeter, confirm each one still receives firmware updates, and retire the ones that do not. An unsupported security appliance is not a security control; it is an unlocked door with a reassuring logo on it.

JUDGE BEHAVIOR, NOT APPEARANCE – AND OPEN THE ENCRYPTED TRAFFIC

When 95.72% of the malware you encounter has been built for exactly one machine, signature matching is structurally outmatched. There is no prior sighting to match against, because there was no prior sighting. Our own detection mix reflects the shift: behavioral and machine-learning engines went from roughly 40% to over 67% of all endpoint detections in a single quarter, and network IntelligentAV detections rose 49% while signature-based Gateway AntiVirus fell 43%. The engines that watch what a file does are carrying an ever-larger share of the load.

So make sure those engines are actually turned on, everywhere. On the network, that means running IntelligentAV and APT Blocker alongside Gateway AntiVirus rather than instead of it. On devices running all three, 28% of malware required proactive detection to catch at all. On endpoints, it means EDR-grade behavioral detection rather than legacy antivirus. Every tier of WatchGuard's refreshed Endpoint Security portfolio now includes AI-powered EDR by default, and if you are running an older configuration, this is the half to revisit it. Where you lack the staff to chase behavioral alerts, a managed service such as WatchGuard MDR closes the gap between a detection firing and someone acting on it.

None of that helps with traffic you never look at, which brings us to the most stubborn statistic in this report. Ninety-five percent of the malware our Fireboxes caught arrived over TLS, and only 20% of deployed Fireboxes inspect encrypted traffic. On devices that do inspect it, the zero-day share of malware jumps from 28% to 36% – the evasive threats are deliberately choosing the channel most of us are not watching. HTTPS inspection is a free feature of every Firebox, regardless of license. It does require real work, like certificate distribution and policy tuning. But an inspection-free perimeter in 2026 means you are seeing, at best, one malware delivery in twenty.

GUARD THE KEYS AND THE UTILITY CLOSET

The most telling finding in our Q1 threat hunting data is that credential access was the single most-alerted ATT&CK tactic, led by stealing passwords straight out of browsers, with credentials in files, cached credentials, and credentials in the registry close behind. Trust control evasion was the second most-invoked rule. Add the FortiBleed credential campaign visible in our Q1 VPN and remote access detections, and the picture is unambiguous: attackers are not battering our controls, they are walking past them holding valid credentials for a trusted account.


Defending the keys starts with phishing-resistant multi-factor authentication (MFA) on every remote access path (VPN, remote desktop, cloud administration, and email) so that a stolen password buys an attacker nothing on its own. WatchGuard AuthPoint is one way to get there, but the specific product matters less than the coverage; a single unprotected VPN portal undoes the rest. Pair that with browser hygiene, since browser-stored passwords were the single most-abused credential source we observed. Discourage saving corporate credentials in browsers, deploy a managed password manager instead, and treat any endpoint alert involving browser credential stores as a probable active intrusion, not a nuisance detection.



Then secure the utility closet, because once inside, attackers work with your own tools. Windows LotL binaries keep taking share as an endpoint attack vector, with Q2 detections. Restrict unnecessary PowerShell functionality and enable script block logging, keep AMSI integrations active and alert loudly when something tries to disable them, monitor for unusual parent-child process relationships and unexpected use of rarely touched system binaries, and control credential-sensitive utilities that touch LSASS. On Linux hosts, add integrity monitoring for /etc/ld.so.preload and system libraries, which is exactly how the process-hiding rootkit in our malware section conceals a miner in plain sight. All of this is also your best ransomware control: extortion groups multiplied this half, and these actors typically arrive by way of stolen credentials and trusted binaries long before an encryptor ever runs.

CONGRATULATIONS.

You've reached the end of our bi-annual, 2026 Global Threat Report. Be sure to come back early next year to keep up with the latest changes in the threat landscape. 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

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.

