







AI SOC Report 2026





Foreword

In 2025, Intezer AI SOC analyzed more than **25 million security alerts** across live enterprise SOC environments. This report is grounded in the triage, investigation, and resolution of every one of those alerts, at a forensic level, across Intezer's global customer base.

With this massive scale of alert volume analyzed, reality is revealed about how attacks actually unfold, where defenses fail, and how SOC's spend their time under pressure.

The big reset of "acceptable security risk"

For years, security leaders have made tradeoffs in the face of an unsolvable equation where alert volumes rapidly grow and human teams do not. Whether operated in-house or through an MDR, most SOC's are forced to triage aggressively, ignoring and auto-closing a majority of alerts, particularly those that are low-severity or informational.

Surveys consistently show that over 60% of alerts are never reviewed, simply because it is mathematically impossible for human analysts to investigate everything. This reality has become the accepted status quo, and understandably so.

But the emergence of AI-driven SOC platforms that can perform high-fidelity analysis at scale, completely resets the definition of "acceptable risk". It's now an imperative to **reexamine long-standing risk tolerance**, especially when data shows that real threats originate in places we historically had no choice but to ignore.

Nearly 1% of [all incidents](#) are traced back to alerts classified at the lowest severity levels.

Real threats hiding in low-severity alerts

One of the most consequential findings in this report is how frequently genuine threats originate from alerts initially labeled as low-severity or informational. Across endpoint, cloud, identity, network, and phishing categories, nearly 1% of all incidents are traced back to alerts classified at the lowest severity levels. At scale, this is far from trivial.

With enterprises generating an average of 450,000 alerts annually, this equates to roughly 54 real threats per year, or about one per week, slipping by unnoticed.

This equates to roughly 54 real threats per year, or about one per week, [slipping by unnoticed](#).



Within endpoint alerts, that number rose to nearly 2% low-severity endpoint alerts representing genuine threats.

In practice, these threats are rarely uncovered by in-house or external SOC teams. Most AI SOC platforms that exclusively rely on AI agents, are also limited to triaging high- and medium-severity alerts. These **“below the radar” alerts** can only be identified through **highly scalable, automated and forensic-grade triage** that can examine behavior, lineage, and intent, for 100% of all ingested alerts.

Noise, stealth, and misjudged risk

Beyond severity misclassification, the data highlights a threat landscape dominated by noise masking meaningful activity. Scanning of public-facing infrastructure generated some of the highest alert volumes, while identity telemetry was flooded with “impossible travel” detections.

In cloud environments, however, the most frequent alerts related to defense evasion and persistence, reflecting attackers’ focus on stealth, long-term access, and abuse of legitimate services through token manipulation and obfuscation.

Taken together, these findings point to a consistent pattern. While alert volume is increasing, severity and risk are often misaligned with meaningful threats frequently hidden among signals that are deprioritized or ignored.

The sections that follow break down these patterns in detail and translate them into operational lessons drawn from a year of large-scale forensic triage. This is the State of the SOC as it exists in practice, shaped by millions of alerts, real adversaries, and the constraints security teams face every day.

As always, Intezer’s AI SOC platform along with our experienced security analysts and researchers, stand ready to deliver stronger security outcomes and real risk reduction that scales with your business.

Severity and risk are often misaligned with meaningful threats frequently **hidden among signals that are deprioritized or ignored.**


Itai Tevet

Intezer CEO and Co-Founder


25,000,000

alerts analyzed by Intezer

Over **60%** of alerts are ignored by SOC and MDR teams, especially low-severity ones



nearly **1%** of *real threats* come from low-severity alerts



What this means for enterprise security?

54 real threats per year (~1 per week)

are never investigated by SOC or MDR

For organizations with **450K alerts per year**





Table of contents

	Data & methodology	07
	Bottom line up front (BLUF)	08
Endpoint	What 9.25M alerts revealed about triage shortcomings	10
	Malware deep dive	11
	Windows	12
	Stealers	12
	Ransomware	13
	Linux	14
	Installers	15
	Weaponized software	15
Network	Noise masking real threats, exposure risks, and emerging abuse	16
	Scanners	16
	Code Sandboxes	18
	I see plain-text passwords	20
	Internal networks	20
	External networks	21
Phishing	Scale, signal, and evolving techniques	22
	User reporting creates volume, but not always signal	22
	Attackers are shifting focus from endpoint to users' browser	22
	Brand impersonation and callback scams dominate	23
	Brand impersonation breakdown	23
	Callback Scams	24
	Phishing infrastructure is becoming harder to detect	26
	New phishing threats and attack vectors	26




Table of contents

Identity	High-value signals, high false-positive volume	29
	Most identity alerts are benign	29
Impossible travel	High volume, low risk	30
Cloud	Persistence, evasion, and configuration drift	31
	Not the best practices	32
About Intezer		33
Appendix	The Rise and Fall of CrazyHunter	34



Data and methodology

This report is based on a comprehensive analysis of threat activity observed across Intezer's global customer base throughout 2025. To identify emerging attack patterns, evasion techniques, and security weaknesses, a broad and diverse dataset was examined.



All data used in this study were processed in accordance with strict privacy and security standards.

Intezer's research methodology is designed to ensure no sensitive, personal, or customer-identifying information is accessed, retained, or included in the analysis.

Findings are aggregated and anonymized, focusing solely on technical threat indicators and behavioral trends. This enables Intezer to deliver accurate, large-scale threat intelligence insights while maintaining the confidentiality and trust of our customers.




Bottom line up front (BLUF)

Intezer AI SOC overall stats

- Less than 2% of all the alerts were escalated by Intezer to a human.
- 98% verdict accuracy.
- Sub-minute median triage time.
- Across all alert types, 0.6% of escalations came from low-severity or informational events.

Triaged alert type breakdown



Endpoint stats

- **2% of escalations came from low-severity or informational endpoint alerts.**
- 180 million files analyzed. This is across all major operating systems: Windows, Linux, MacOS, iOS, and Android.
- Over 82,000 forensic investigations performed on endpoints, including live memory scans. In over 1.6% of the forensic scans, we found that the endpoint was still compromised even though the EDR reported the threat had been mitigated.
- Threat actors deploying stealers and RATs use increasingly complex, multi-stage loaders to disguise familiar malware strains.
- Akira Ransomware has topped our results as one of the most active ransomware groups in 2025. Ransomware is still being created with flawed cryptography.
- Strains of Mirai were the most prevalent malware identified on Linux hosts.



Network stats

- Over 1M alerts on IP addresses (out of the 7M investigated) were for scanning activity. 88% were blocked at the source by the firewall.
 - Of the remaining 12% that required triage, only 6% were confirmed as malicious.
- **Code sandboxes are being heavily abused to host phishing sites targeting the theft of cryptocurrency.**
- A significant amount of internal network traffic remains unencrypted during transit. Indicates that many companies still rely on perimeter security rather than zero-trust security.

Phishing stats

- Over half a million user-reported phishing emails were analyzed. The most common technique used in phishing emails was brand impersonation, followed by callback scams and credential phishing links.
- Over 3 million phishing and malware-related URLs investigated. Microsoft was the most impersonated brand.
- **Threat actors have increased their use of CAPTCHA to make the detection of phishing pages more difficult.**
- **PayPal's infrastructure is being heavily abused for callback scams.**
- Malicious content is being concealed within trusted file formats to bypass security measures.

Identity stats

- The most common identity alert was for location or session anomalies. Over 36% of identity alerts fall into this category. Most "impossible travel" alerts are false positives, mainly caused by VPNs and mobile phones.

Cloud alert stats

- Most of the cloud-based alerts fall in the defense evasion and persistence category.
- **Amazon S3 is involved in the majority of AWS misconfigurations.** Many legacy features are still being used.




Endpoint: What 9.25M alerts revealed about triage shortcomings


In 2025, Intezer AI SOC investigated 9.25M endpoint alerts, analyzing over 175 million files, combining data from the alerts with Intezer's proprietary endpoint scans to achieve a verdict.

Of these endpoint alerts, **43% were classified as informational or low-severity.**


Of these informational or low-severity alerts, **nearly 2% of them were real incidents.**



Additionally, over half of all endpoint alerts were not automatically mitigated by their endpoint protection solution. **Of these non-mitigated alerts, almost 9% were confirmed as malicious.**



Alerts that required additional analysis, underwent live endpoint scans with Intezer's proprietary endpoint scanner. In **more than half of the confirmed compromised endpoints, the endpoint protection reported the alert as mitigated!**



Note

While it's possible these numbers reflect a vendor issue when it comes to mitigating or correctly labeling threats, there can be many other legitimate factors at play causing this behavior. All that is certain is that SOC teams need to be vigilant in light of these realities.




Malware deep dive

One of Intezer's proprietary investigative tools is an endpoint scanner which performs live memory forensics on a host and identifies malicious code running within a process.

This capability makes our endpoint analysis not only comprehensive but also highly effective in uncovering memory-resident threats, backdoors, and advanced persistent activity that may originate from seemingly low-severity alerts.

Top malware families compromising endpoints (detected in memory)





Windows

Windows remains the most popular operating system for endpoints, accounting for 51% of all endpoints. As such, we also see a wide range of threats targeting these endpoints. In this section, we will highlight the most common types of threats that we saw in the past year.

Breakdown of operating systems


 Windows	51%
 Linux	46%
 macOS	1.5%
 iOS	<1%
 Android	<1%

Breakdown of devices

 Desktop	53%
 Server	46%
 Virtual	<1%
 Phone/tablet	<1%

Stealers

Credential and information stealers remained a prevalent threat throughout the year, with Agent Tesla dominating the landscape, accounting for 39% of all detections. It was followed by Snake Keylogger and Formbook, which together accounted for more than a third of activity, highlighting attackers' continued focus on harvesting sensitive data and credentials for follow-on intrusions or resale. Other notable families, such as Mimikatz, LummaC2, and RedLine Stealer, further illustrate the widespread use of credential theft as an initial access and persistence mechanism across both criminal and state-aligned campaigns.






Ransomware

The ransomware activity observed over the past year shows a clear concentration around a few dominant families with Akira, INC, and LockBit appearing as the most active groups by a significant margin. These operations continued to drive a large portion of global ransomware incidents, reflecting their mature infrastructure, rapid development cycles, and broad targeting across various industries. Beyond these leading families, well-known groups such as Babuk, Lynx, and STOP also remained

active, alongside smaller but persistent clusters like Phobos, AvosLocker, and Rhysida. Even lower-volume families such as CHAOS, Pay2Key, and Trigona demonstrated continued activity, illustrating the diverse and evolving nature of the ransomware ecosystem. The data shows a ransomware landscape dominated by mature RaaS operations while new variants continue to appear, creating a mix of stable, high-impact groups and ongoing experimentation by threat actors.

Common ransomware families






Linux

Mirai remains the most prevalent threat in Linux environments, reflecting the continued exploitation of vulnerable IoT and server infrastructure. Its dominance stems from the widespread use of weak credentials, exposed services, and unpatched devices, enabling the automated propagation of malware at scale. Despite the emergence of ransomware and cryptominer families like Akira and XMRig, none rival Mirai's persistence or reach. The malware's open-source availability and frequent

adaptation for different architectures make it a versatile tool for both low-level and organized threat actors. Its continued activity underscores ongoing weaknesses in IoT security and highlights the need for stronger authentication, segmentation, and proactive monitoring across Linux-based systems.

We also observe ransomware families targeting Linux systems alongside Windows environments, with Akira and INC remaining the most dominant across Linux machines.


Malware threats in Linux





Installers

The pie chart here shows the installers most commonly used as the initial delivery mechanism in malware cases observed on our platform. Because many malicious payloads are packaged inside legitimate installer frameworks, the root file often appears benign while the installed components are not. With Intezer's unique sandboxing and genetic analysis data, the installer family was tracked when the resulting installed file was confirmed as malware. The breakdown highlights how frequently specific installers are abused for this purpose.



Weaponized software

- **Newly observed multi-stage loader called PNGPlug delivers the well-known remote access trojan, ValleyRAT**

Our research team uncovered a campaign that utilized a newly observed multi-stage loader called PNGPlug to deliver the well-known remote access trojan, ValleyRAT. The attack began with a phishing page that distributed a malicious MSI installer posing as legitimate software. Once executed, it unpacked and decrypted additional components, using a large loader DLL that injected payloads hidden within PNG files before launching ValleyRAT. This discovery highlights a growing trend where stealers and RATs use increasingly complex, multi-stage loaders to disguise familiar malware strains. Attackers are focusing on layered execution and loader obfuscation to bypass traditional detections, making behavioral and memory-based analysis critical for identifying threats early in the infection chain. [Read the full report.](#)

- **Zero-day vulnerabilities actively exploited by the XE Group**

threat group In collaboration with Solis, a U.S.-based MSSP, Intezer identified two previously unknown zero-day vulnerabilities that were being actively exploited by the XE Group threat group. The investigation revealed that XE Group had evolved from conducting web-skimming campaigns to leveraging advanced zero-day exploits targeting the VeraCore warehouse management platform. These exploits ([CVE-2024-57968](#) and [CVE-2025-25181](#)) enabled attackers to upload malicious ASPX web shells, execute obfuscated payloads, and deploy in-memory shellcode using reflective PowerShell loaders and Meterpreter, thereby allowing for stealthy persistence and data exfiltration. [Read the full report.](#)

- **The rise and fall of CrazyHunter**

In early 2025, Intezer and Taiwan-based AISHield collaborated on an investigation into a newly emerged ransomware group dubbed CrazyHunter, which launched a wave of coordinated attacks against Taiwan's healthcare, education, and energy sectors. The campaign stood out for its sophistication, despite the group's apparent newness, blending traditional double-extortion tactics with modern evasion techniques. [See appendix for full details.](#)




Network: Noise masking real threats, exposure risks, and emerging abuse

The network layer provides broad visibility into attacker behavior but also generates significant noise. In 2025, Intezer AI SOC triaged alerts involving over 7 million IP addresses, 3 million domains, and 3 million URLs, with contextual analysis enabling us to distinguish routine activity from meaningful risk.

Scanners

Web scanners account for a significant portion of the alerts triaged by Intezer AI SOC with over one million alerts this year alone. This volume reflects the extensive number of public-facing servers we monitor and the constant probing activity that targets exposed infrastructure across the internet.

Most scanner activity we observed was automatically mitigated by firewall products with many cases traced to benign scanners such as Shodan and Censys. **Of the 12% of alerts not mitigated by firewalls, 6% were confirmed as genuine threats.**





The following table highlights the ten most active malicious scanners observed throughout the year, ranked by the number of alerts in which they appeared. This view illustrates the persistence of automated reconnaissance activity across monitored environments

IP Address	Number of Alerts	Type
79.124.62[.]122	4,462	Port Scanner
45.224.42[.]19	1,095	SSH Brute Force
46.17.96[.]38	656	Port Scanner
185.44.76[.]10	564	Port Scanner
45.142.193[.]107	519	Port Scanner
139.144.52[.]241	478	Vulnerability Scanner
195.140.213[.]28	474	Port Scanner
141.98.80[.]118	438	Password Brute Force
172.245.163[.]134	417	Port Scanner
45.134.26[.]32	414	Port Scanner

Any system exposed to the internet will inevitably attract scanner activity, whether benign or malicious. The key to effective triage lies in context: combining alert data with threat intelligence and insight into the affected device or account to determine whether an actual breach has occurred.




Code Sandboxes

While reviewing network alerts from XDR and SIEM sources, we noticed recurring references to well-known code sandboxes. Initially, the findings appeared benign. These domains are legitimate platforms used by developers to share and test code, and they maintain a solid reputation across most threat intelligence feeds. Yet despite that, Intezer AI SOC was returning malicious verdicts.

We examined the alerts and found that although the parent domains were reputable, the individual projects hosted on them often served phishing pages. In several cases, the pages were live credential-harvesting sites using the sandbox as free, trusted hosting.







Code sandboxes are attractive to attackers because they offer immediate deployment, a trusted domain, and no infrastructure cost. The content they host runs as intended, which allows a malicious page to operate without technical restrictions. These characteristics make the environments ideal for short-lived phishing campaigns that benefit from built-in credibility.

The platform most frequently abused in our data was Vercel. It appeared across a wide range of phishing themes, including social media, email, banking, streaming, and gambling. One example involved a link sent by email that led to a Vercel page imitating Adobe Cloud, prompting users to sign in to view a fake document.





A large number of incidents also involved cryptocurrency scams. We found sites asking victims to enter recovery phrases or credentials for platforms such as MetaMask, Coinbase, Binance, and Uniswap. Several targeted NFT marketplaces, such as OpenSea, are indicating that attackers are expanding their focus to the broader Web3 ecosystem.



Vercel was not the only platform misused. Similar activity appeared on CodePen, JSitor, and JSBin, each hosting copies of phishing kits or redirect pages; all of which have been very reactive at taking down malicious content. The pattern suggests a continuing shift toward the abuse of legitimate developer services as disposable infrastructure for credential theft.




I see plain-text passwords

Plain-text password alerts usually originate from products that inspect network traffic, such as DLP systems, XDR engines with network visibility, or SIEM pipelines that analyze packet metadata.

When we reviewed our own data, we found roughly a five-to-one ratio of internal detections vs. external ones (In this analysis, “internal” traffic refers to communication where both the source and destination use internal (private) IP addresses, such as RFC1918 ranges (e.g., 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16). “External” traffic refers to communication where at least one side uses a public IP address, even if both systems belong to the same organization).

This ratio reflects the difference in how often internal and external traffic is encrypted as these alerts are raised when authentication information is sent without transport-layer encryption.



Internal networks

Within internal networks, most plaintext-password detections were linked to directory services and other authentication mechanisms that still rely on unencrypted LDAP. LDAP is widely used within organizations for user lookups, access control, and interactions with shared resources. When it is not configured with LDAPS or StartTLS, it transmits credentials in cleartext.

This made it the most common source of internal alerts. We also saw cleartext passwords in internal web services, file-transfer utilities, and database connections. These internal exposures were consistent throughout the year, reflecting long-standing configurations and legacy applications rather than short-term mistakes.

Rank	Top Apps with Leaked Passwords (Internal)
1	Lightweight Directory Access Protocol (LDAP)
2	Hypertext Transfer Protocol (HTTP)
3	File Transfer Protocol (FTP)
4	Databases (TNS/Postgres/TDS)
5	Telnet



External networks

External plaintext-password alerts showed a very different pattern. The most frequent cases involved database protocols, unencrypted file transfer traffic, and HTTP sessions without TLS encryption. In many of these alerts, both endpoints actually belonged to the same organization but communicated with each other through the public IP address of the service rather than the internal one.

This behavior is common in cloud or containerised deployments, where services are exposed externally through load balancers or NodePort-style configurations and then inadvertently used by internal components. Many external alerts also appeared in spikes and then stopped, which suggests temporary misconfigurations that were corrected quickly after detection.

Taken together, **the datasets show two distinct behaviors. Internal plaintext-password traffic tends to come from established systems that have not yet been modernized to use encrypted transport. External plaintext-password traffic is usually the result of deployment or routing mistakes in more dynamic environments, and is often corrected once noticed.**

These exposures relate directly to the OWASP Top Ten categories of [Cryptographic Failures](#) and [Security Misconfiguration](#), which highlight the risks created when sensitive data is transmitted without encryption or when services are unintentionally exposed or configured without proper safeguards.


Rank	Top Apps with Leaked Passwords (External)
1	Databases (TNS/Postgres/TDS)
2	File Transfer Protocol (FTP)
3	Hypertext Transfer Protocol (HTTP)
4	Simple Mail Transfer Protocol (SMTP)
5	Apache Subversion (SVN)



Phishing: Scale, signal, and evolving techniques

Phishing remains a persistent challenge for security teams, not because of detection gaps alone, but due to volume, user behavior, and rapidly evolving attacker techniques.

In 2025, Intezer AI SOC processed **550,000 phishing emails reported by end-users** that bypassed standard email security controls. A bit less than **8% were confirmed as malicious**, highlighting both the value and the operational cost of user-driven reporting.



User reporting creates volume, but not always signal

While security awareness training encourages reporting suspicious emails, a small number of overzealous users can generate disproportionate reporting volume, creating noise that can overwhelm security teams and slow response to genuine threats.

Attackers are shifting focus from endpoint to users' browser

Less than 6% of reported phishing emails had an attachment and around 30% used links. This means that artifact-based analysis is less efficient at triaging phishing emails, as most of them are language-based and require a language model analysis process. The higher prevalence of links over attachments indicates that threat actors are shifting their focus from gaining a foothold on the endpoint, which is well-protected, to instead staying within the user's web browser, where today's security products have limited visibility.




Brand impersonation and callback scams dominate

Among confirmed malicious phishing emails:

- Brand impersonation was the most common technique.
- Callback scams ranked second, often abusing legitimate services to appear credible.


Taken together with brand impersonations, **more than a third of malicious phishing emails are abusing users' trust** in well-known brands.



Brand impersonation breakdown

Microsoft and DocuSign accounted for nearly 85% of brand impersonation cases, underscoring how attackers exploit trusted brands to increase success rates. While DocuSign essentially provides one product, Microsoft offers many products and services that are being impersonated by threat actors.


Right behind impersonations of someone within the victim's organization (6%), **PayPal, accounted for almost 4% of all brand impersonation phishing emails**. While PayPal was more popular in callback scams since those emails were sent by PayPal, we don't classify them as impersonation emails.





For the phishing URLs that performed brand impersonation, we see a similar trend.

Microsoft was the most impersonated brand representing nearly one-third of impersonation URLs. Generic login pages represent almost 15% and Cloudflare almost 14% of URL impersonations.



Callback Scams

At the beginning of 2025, we started to see a **callback scam that abused both PayPal's and Microsoft's infrastructure**. Fortinet described how the attack is performed in their January blog [post](#). In summary, the threat actor registers for a test domain on Outlook 365 and creates a distribution list.

In PayPal, the threat actor creates a payment request that is sent to the distribution list. This allows the threat actor to send one legitimate email from PayPal to multiple victims that passes all the standard email security checks.

This phishing email has two attacks. First, as Fortinet described in the blog, if the victim clicks the Pay Now link, they are redirected to PayPal's website. PayPal wants to make it easy for users; it will link the receiving email address to your PayPal account if you choose to log in. Because the attacker sent the phishing to their controlled email address, this is the email address that will be added to the victim's PayPal account.

PayPal

julie garfield has requested

\$ 299.99 USD

We'll link **admin104@menloparks.org** to your PayPal account when you log in.

"Fund Transfer Request Approved and Payment D..."

Email address

Password

Log In to Pay

[Having trouble logging in?](#)

Pay Without a PayPal Account



The second attack is in the email. In the note section of the payment request, the attacker places a message that encourages the victim to call a specific phone number. The message is designed to appear as if it is from PayPal and requests the recipient of the email to call the provided phone number if they don't recognize the payment or have any questions.

Throughout the year, we have seen threat actors evolving this attack. The recent phishing emails are also using G-Suite in the delivery chain. The phishing email in the figure was sent from PayPal to a mailbox hosted on G-Suite that bounced the email to an email address hosted on Outlook 365. The Outlook 365 mailbox is used to forward the email to a distribution list that contains the victim's email address. You can also see in the screenshot that Unicode homoglyphs are being used to make signature-based detection harder.

Fromservice@paypal.com

10/17/25, 16:37

Toadmin731@sunnyhomespain.it.com

SubjectYou've got a money request

ReplyReply AllForwardArchiveJunkDeleteMore

To protect your privacy, Thunderbird has blocked remote content in this message.

Preferences

ROYAL FLUSH

ENTERTAINMENT sent you a

money request

Payment request details

Amount requested

\$299.00 USD

Note from ROYAL FLUSH ENTERTAINMENT:

Fund Transfer Request Approved and Payment Deducted. FOR ANY INQUIRIES, CONTACT: +1 (888) 404-1720

Transaction ID

U-5JJ66757BK416974F

Transaction date


October 17, 2025

Pay Now



Phishing infrastructure is becoming harder to detect

Threat actors have increased their use of CAPTCHA to make the detection of phishing pages more difficult and apparently prefer Cloudflare's Turnstile over Google's reCAPTCHA. The majority of URLs that use Google CAPTCHA could be classified as safe, while the opposite is true for sites using Cloudflare Turnstile.



New phishing threats and attack vectors

As a benefit of analyzing a large amount of phishing emails and URLs, Intezer uncovered a striking shift in attacker strategy, from simple deception to structural and contextual evasion.

These campaigns increasingly conceal malicious content within trusted file formats, cloud services, and dynamic web components, which are designed to evade modern email and endpoint defenses.



- MHT files embedded in OpenXML documents. DOCX attachments conceal archived HTML content or QR codes leading to credential-harvesting portals.

These techniques highlight a growing preference among attackers for non-traditional delivery vectors that exploit file format flexibility, embedded scripting, and encoding, sometimes in combination with the trust associated with legitimate file-sharing platforms.

By manipulating the structure of formats like SVG, PDF, and OpenXML, attackers conceal malicious content in layers that most scanners overlook, allowing phishing payloads to penetrate even well-defended environments.





Identity: High-value signals, high false-positive volume

Identity remains one of the most targeted and valuable control planes in modern environments. In 2025, Intezer AI SOC analyzed millions of identity-related alerts from major identity providers, including Okta, JumpCloud, and Microsoft Entra ID.

These alerts are primarily related to suspicious authentication activity, such as brute-force attempts, failed multi-factor authentication (MFA) challenges, and sign-ins from unusual geographic locations or unfamiliar devices.

The majority of identity alerts were driven by location-based anomalies (37%) and login failures (22%), followed by MFA issues and brute-force attempts. Alerts also covered policy changes, privilege modifications, and access enforcement events.


Rank	Percent
Location / Geo / Session Anomalies	36.70%
Login Failures	22.00%
Brute Force / Password Attacks	8.20%
Account Lockouts	6.20%
MFA Issues	11.90%
Privilege & Group Changes	3.90%
Kerberos / Cert Thumbprint Activity	2.50%
Access Policy Blocks	2.60%
Other	6.00%

Most identity alerts are benign

To better understand the nature of identity-related activity, we analyzed output from our AI classifier, which integrates directly with customers' identity providers and evaluates recent authentication events. The classifier examines login behavior, contextual signals, and historical patterns to identify unusual or risky activity. The breakdown below shows how these events were ultimately classified.

From our data, we can see that most identity alerts are false positives. AI is good at recognising normal variability. The real threats show up only when behavior sharply deviates across IP, geography, and timing.

Login verdicts





Impossible travel: High volume, low risk

Impossible-travel alerts, common features of identity and DLP products, were frequent but rarely malicious. **Only ~2% were confirmed as real compromises.**

Here are some of the causes of the high false-positive rates:

VPN Usage

A large portion of false positives can be attributed to VPN usage. Both commercial VPN services and internally managed VPN gateways contribute to this pattern. When a user switches between their normal connection and a VPN endpoint, the identity provider interprets the traffic as coming from two distant locations in rapid succession. Roughly 30% of all impossible-travel alerts involved some form of VPN activity. The most popular commercial VPNs observed are NordVPN, Proton, SurfShark, Private Internet Access (PIA) and ExpressVPN.

To triage these cases, we rely on tuning informed by threat intelligence, customer feedback, and observed patterns that differentiate legitimate VPN usage from suspicious activity. iCloud Private Relay also appears frequently and links directly to mobile behavior.

Mobile phones

Mobile phones are another major driver of false positives. A typical case involves a user logged in on a workstation, where the traffic originates from the organization's normal geographic region.

The same user then checks email or a messaging application on their phone while using mobile data. Since mobile operators often route traffic through data centres far from the user's physical location, the identity provider flags the login as a separate, distant location. This behavior alone accounts for a significant share of impossible-travel alerts.

Security tools

We also see a pattern where security tools unintentionally trigger alerts against each other. "Zero Trust" solutions, secure web gateways, and some email-security products route or inspect traffic through their own infrastructure. This causes authentication events to originate from cloud locations associated with those services rather than the user's true location. Identity platforms, SIEMs, and DLP systems interpret these as unexpected logins and generate impossible-travel alerts. These cases emphasize the recurring theme of overlapping security products creating noise or misleading signals for one another.



Cloud: Persistence, evasion, and configuration drift


Cloud environments continue to evolve rapidly, introducing new forms of telemetry, behavior, and detection opportunities. Our AI-SOC processes a wide variety of cloud alerts across multiple platforms, including AWS, Azure, and GCP, as well as events generated by SIEM systems and native detection frameworks such as AWS GuardDuty.

Cloud alerts differ from traditional network or endpoint detections in that they often represent behavioural signals rather than discrete malicious artifacts.

The cloud alert data shows a strong concentration of activity around Defense Evasion and Persistence, which together account for the majority of observed TTPs.

These patterns suggest that most threats in cloud environments primarily focus on maintaining long-term access while evading detection, often through techniques such as obfuscation, token manipulation, and the misuse of legitimate cloud features. Discovery, Credential Access, and Resource Development follow as common tactics, reflecting attackers' efforts to explore the environment, harvest credentials, and prepare infrastructure for further operations. In contrast, high-impact behaviors such as Lateral Movement and Privilege Escalation appear far less frequently, suggesting that early-stage stealth and identity-driven techniques dominate cloud intrusions. This distribution highlights the importance of organizations enhancing visibility into identity misuse, covert persistence mechanisms, and subtle evasion behaviors within their cloud environments.

Cloud alerts: seen TTPS and tactics






Not the best practices

Cloud posture tools such as AWS Security Hub continuously evaluate an environment against established security controls.

These controls represent recommended safeguards for ensuring the confidentiality, integrity, and availability of cloud resources. Most of these violations fall into the low or informational categories, which is expected for posture-centric findings where misconfigurations are more common than immediate threats.



Across all services, S3 accounts for the largest share of control violations. When the alerts are grouped by service, roughly 70% of all findings relate to S3 security hardening. The most frequently violated S3 controls are those requiring SSL for bucket access, enforcing server access logging, and restricting bucket policies to prevent cross-account access.

Taken together, these patterns suggest a recurring theme: many organizations still operate S3 buckets with default or legacy settings that pre-date modern security expectations. Missing access logs reduces visibility, weak cross-account policies increase the chance of accidental exposure, and a lack of enforced encryption in transit leaves room for downgrade attacks or insecure client behavior. These issues are typically configuration oversights rather than malicious activity, but they represent the type of long-lived weaknesses that attackers routinely exploit when they gain a foothold elsewhere in an environment.

Rank	Control Code	Control Description
1	S3.12	ACLs should not be used to manage user access to S3 general purpose buckets
2	S3.9	S3 general purpose buckets should have server access logging enabled
3	S3.6	S3 general purpose bucket policies should restrict access to other AWS accounts

EC2 findings form the next major cluster. The most common violations involve EC2 instances using multiple network interfaces, the presence of instances with public IPv4 addresses, and workloads running on paravirtual or otherwise outdated instance types. These issues point to operational debt and legacy design patterns persisting in production. Multiple ENIs are often linked to older architectures, complex routing setups, or automation that attaches extra interfaces by default, and they can expand the attack surface and complicate segmentation. Public IPv4 assignments similarly increase exposure to external probing. The repeated appearance of deprecated instance families suggests that some workloads remain on historical templates or provisioning pipelines that have not been refreshed to modern defaults.




While these controls are often classified as low-severity, the associated risk compounds when combined with other cloud misconfigurations, particularly in environments without strong network segmentation.

Rank	Control Code	Control Description
1	EC2.17	EC2 instances should not use multiple ENIs
2	EC2.24	EC2 paravirtual instance types should not be used
3	EC2.9	EC2 instances should not have a public IPv4 address

When comparing S3 and EC2 together, a clear pattern emerges. S3 violations tend to reflect broad, environment-wide defaults that were never fully aligned with AWS recommendations. EC2 violations, on the other hand, often reflect application- or team-specific decisions, legacy infrastructure, or under-maintained automation pipelines. S3 issues cluster around access control and encryption; EC2 issues cluster around exposure and outdated configuration. Both categories highlight the gap between theoretical cloud-security best practices and the reality of fast-moving infrastructure that accumulates configuration drift over time.

Control Violations per AWS Product



About Intezer

Intezer AI SOC delivers 24/7, forensic-grade cyber alert triage across 100% of alerts, with only 2% escalated for human review, dramatically accelerating incident response. Powered by ForensicAI™, Intezer specializes in deep forensic investigation to deliver unmatched accuracy and speed, significantly reducing cyber risk and enabling security teams to operate effectively without reliance on outsourced services. Intezer is trusted by global enterprises including NVIDIA, MGM Resorts, Equifax, Salesforce, and Ferguson.




Appendix

The Rise and Fall of CrazyHunter

In early 2025, Intezer and Taiwan-based AIShield collaborated on an investigation into a newly emerged ransomware group dubbed CrazyHunter, which launched a wave of coordinated attacks against Taiwan's healthcare, education, and energy sectors. The campaign stood out for its sophistication, despite the group's apparent newness, blending traditional double-extortion tactics with modern evasion techniques.

The joint investigation combined AIShield's local telemetry with Intezer's deep malware analysis expertise to deliver a full technical breakdown of CrazyHunter's toolset and infrastructure. The group's operations revealed the use of vulnerable drivers for privilege escalation, Active Directory Group Policy abuse for lateral movement, and a Tor-based leak site and Telegram channel for ransom negotiations and public exposure.

Among the tools analyzed were the ZammOcide AV-killer, which exploits CVE-2018-6606, a loader called "Syscall Phantom" that bypasses EDR hooks, and GoStealthFile, a Golang utility for file exfiltration and remote access. Their main payload, "Hunter Ransomware," was derived from an open-source project but reconfigured with their own branding and modifications. While its encryption implementation was not the most robust, the campaign's orchestration demonstrated real operational maturity.



DARK WEB EXTORTION SITE USED BY CRAZYHUNTER



CrazyHunters Arsenal

Tool	Category	Description
Hunter Ransomware (Prince)	Ransomware	Open-source Prince-based encryptor rebranded as "Hunter"; encrypts files and drops ransom note.
ZammOcide (C)	AV-killer	Exploits vulnerable anti-malware driver (CVE-2018-6606) to terminate antivirus processes.
Golang ZammOcide	AV-killer	Golang port of ZammOcide used similarly to disable endpoint protection.
Syscall Phantom	Loader / Evasion	Uses manual syscalls to bypass EDR hooks and inject shellcode in memory.
GoStealthFile	Exfiltration / Cleanup	Golang utility for selective file deletion or serving files via embedded HTTP server.
SharpGPOAbuse	Lateral movement	Open source .NET toolset to manipulate Group Policy Objects and schedule tasks for persistence.
Batch Orchestrater	Orchestration	Staged bootstrap script coordinating AV killers, loaders, and ransomware execution.
Cobalt Strike Beacon	C2 / Beacon	Cobalt Strike HTTPS beacon used for command-and-control and post-exploitation.
Rex::Powershell (modified)	Loader script	PowerShell loader performing XOR decryption and in-memory beacon loading.

Taiwanese authorities have identified the operator behind the CrazyHunter ransomware as a man surnamed Luo (羅) from Zhejiang, China. Using IP data, malware code similarities, and behavioral analysis, investigators linked him to attacks that compromised over 500 systems and exposed about 32.5 GB of patient data affecting more than 16 million people.

The case has been handed to the Taipei District Prosecutors' Office, which issued a wanted notice for Luo on charges including computer misuse, data protection violations, and extortion. This confirmation marks a rare and significant instance of clear attribution in a ransomware investigation.



