Wiz Research Team

ZERO-DAYS | CLOUD MISCONFIGS | RESEARCH | THREAT HUNTING | INCIDENT RESPONSE | OH NO MOMENTS

WIZ

The Crying Out Cloud

CHRONICLES

Latest Vulnerabilities & Threats Uncovered by Wiz Research

Issue Vol.1 July-August 2025 $4.00

Ingress-NGINX Kubernetes Vulnerabilities Pose RCE and Info Disclosure Risks

Researchers have uncovered several critical vulnerabilities in the popular ingress-nginx controller for Kubernetes, potentially affecting thousands of exposed clusters. These issues, which include remote code execution (RCE) and information disclosure risks, stem from insecure handling of annotations and configuration values. The ingress-nginx controller, widely used for routing traffic to Kubernetes applications, now poses an unintentional security risk for many deployments.

Two CVEs have been assigned to the most severe vulnerabilities, with exploitation allowing attackers to manipulate ingress configurations in ways that execute arbitrary code or exfiltrate sensitive environment variables. The vulnerabilities are especially concerning given the controller's deep integration into production workloads, which could provide threat actors a direct pathway into cloud-native infrastructure.

Hillai Ben-Sasson, Wiz Research

To address the issue, the ingress-nginx project maintainers have released patched versions and security advisories urging users to upgrade immediately. Organizations should review ingress resources for potentially unsafe annotations, limit the use of external inputs, and apply least-privilege principles to Kubernetes service accounts.

GitHub Action Supply Chain Attack Hits Popular CI Workflow

A widespread supply chain attack has been discovered in the GitHub Action ecosystem, targeting the popular tj-actions/changed-files workflow used in thousands of CI/CD pipelines.

Tracked as CVE-2025-30066, this vulnerability was exploited via a malicious fork of the action, which was pushed to projects through automated pull requests. Unsuspecting developers merged these updates, unknowingly introducing malicious logic into their build workflows.

The attack enabled the exfiltration of secrets and manipulation of CI workflows, posing a serious threat to both open-source and enterprise development environments. Because GitHub Actions operate with high trust, any compromised action can quickly lead to widespread impact. This incident underscores the ongoing risks of software supply chain attacks and the need for careful dependency management.

> "*Unsuspecting developers merged these updates, unknowingly introducing malicious logic into their build workflows.*"

GitHub has responded by removing the malicious fork and revoking related tokens, while encouraging developers to verify the source of third-party actions. Users are urged to audit workflows, pin trusted versions, and use actions from verified publishers.

Exposed DeepSeek Database Leak: AI Data at Risk

Wiz Threat Research revealed a massive data leak involving DeepSeek, an AI research initiative. The exposed database contained a vast repository of user queries and AI-generated responses, making it a potential goldmine for cybercriminals. This unprotected data could be exploited for social engineering, spear-phishing campaigns, and AI-driven misinformation. The breach serves as a stark reminder of the risks associated with storing sensitive AI-generated interactions without proper security measures.

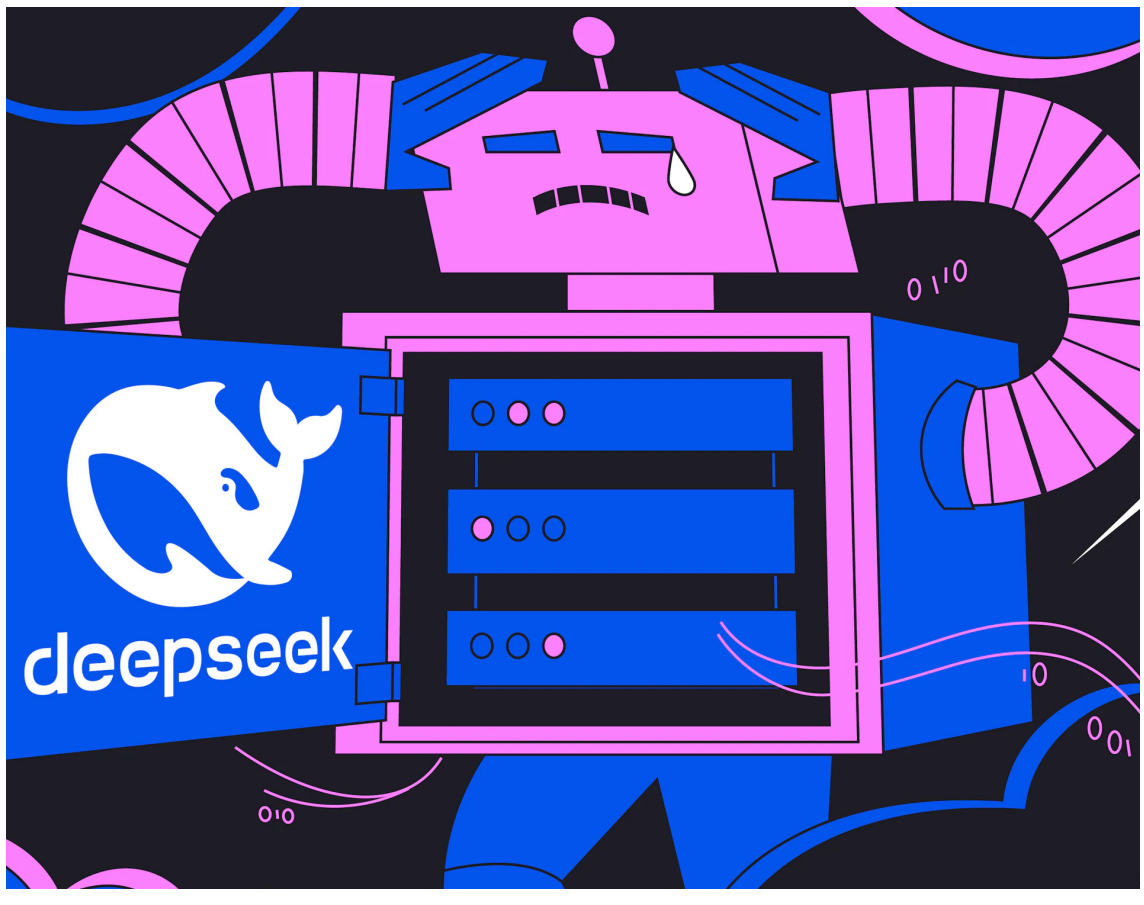

The lack of access controls meant that anyone with an internet connection could retrieve and analyze the stored information. This poses significant threats, as malicious actors can leverage AI responses to refine phishing emails, impersonate individuals, or craft highly convincing fake interactions. With AI playing an increasingly important role in decision-making and communication, breaches of this nature have far-reaching consequences beyond simple data leaks.

To prevent such incidents, organizations working with AI-generated data must implement stringent security measures, including robust authentication protocols, encryption, and continuous monitoring for unauthorized access. Regulatory frameworks should also be updated to ensure AI-related data is properly secured, preventing the potential misuse of valuable and sensitive information.

CVE-2024-50603: Exploited in the Wild

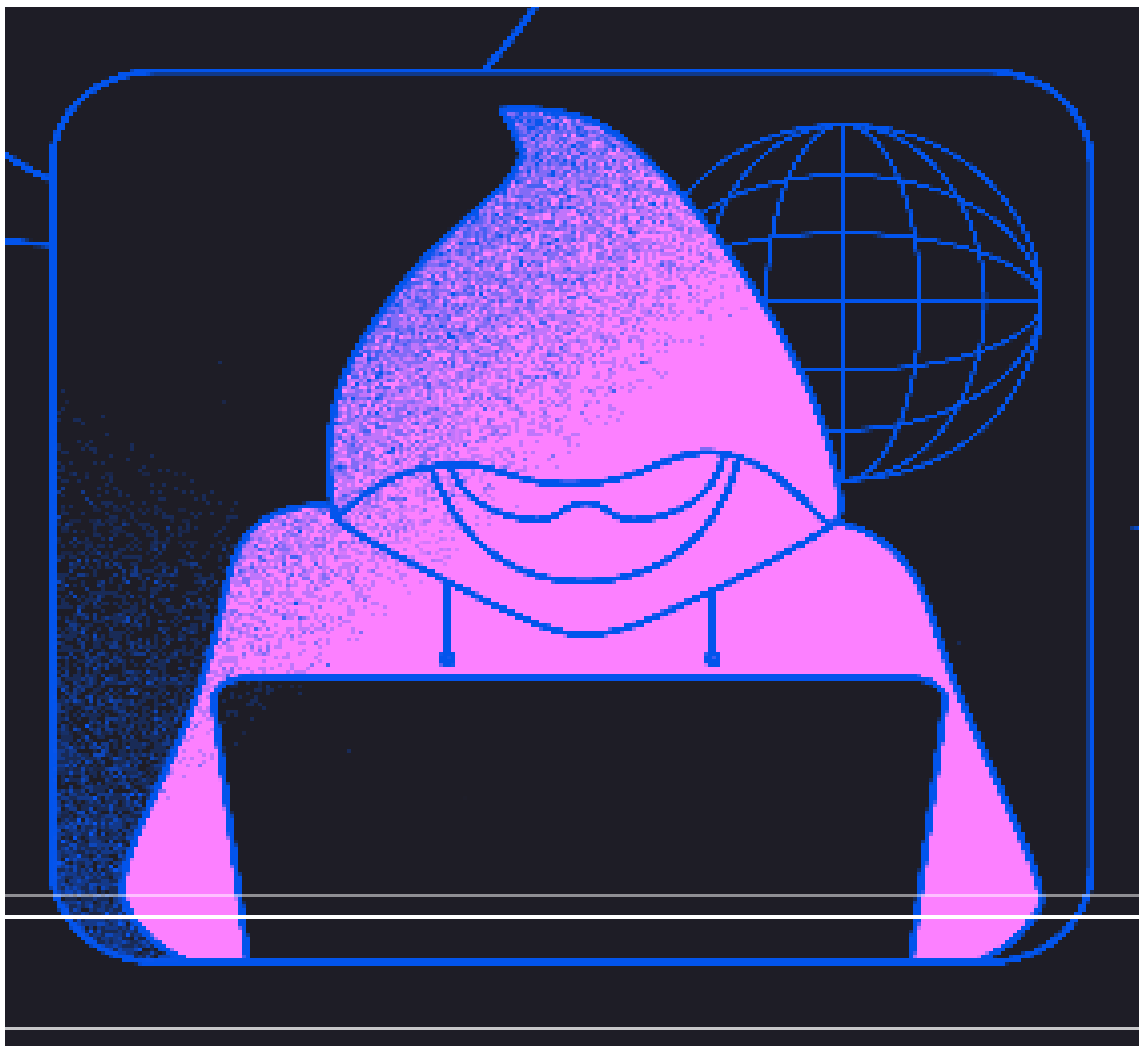

A critical vulnerability, CVE-2024-50603, has been identified in Aviatrix VPN software, and reports confirm that it is currently being exploited in the wild. The flaw, which allows remote code execution, enables attackers to gain full control over affected systems.

Wiz Threat Research observed real-world exploitation of this vulnerability, with attackers using it to establish persistent access, execute arbitrary commands, and deploy malware. The nature of the vulnerability makes unpatched systems an easy target, significantly increasing the urgency for organizations to apply updates immediately. Failure to do so could result in data breaches, service disruptions, or worse—full system compromise. To mitigate the risks associated with CVE-2024-50603, organizations should update to the latest patched version of Aviatrix immediately. Additionally, ***network segmentation,*** multi-factor authentication, and continuous monitoring of VPN access logs can help minimize the potential damage caused by an attack.

"Wiz Threat Research observed real-world exploitation of this vulnerability, with attackers using it to establish persistent access, execute arbitrary commands, and deploy malware."

DIICOT THREAT GROUP LAUNCHES MALWARE CAMPAIGN

Wiz Threat Research observed the cybercriminal organization DIICOT has launched a highly sophisticated malware campaign aimed at enterprises and governmental institutions worldwide. Using an intricate phishing campaign, the attackers successfully compromise systems by embedding malware into seemingly legitimate email communications. The malware then gains access to critical data, including credentials, financial information, and classified communications.

Once installed, the malware uses advanced techniques to avoid detection, such as encryption, obfuscation, and leveraging trusted cloud services for command and control operations. This enables attackers to persist within a victim's environment for extended periods, gathering intelligence and executing secondary payloads when necessary. The adaptability of DIICOT's campaign suggests that the group is well-funded and experienced, posing a severe threat to cybersecurity.

To defend against such threats, organizations must enhance email filtering, conduct regular security awareness training, and deploy endpoint detection solutions capable of identifying suspicious activity.

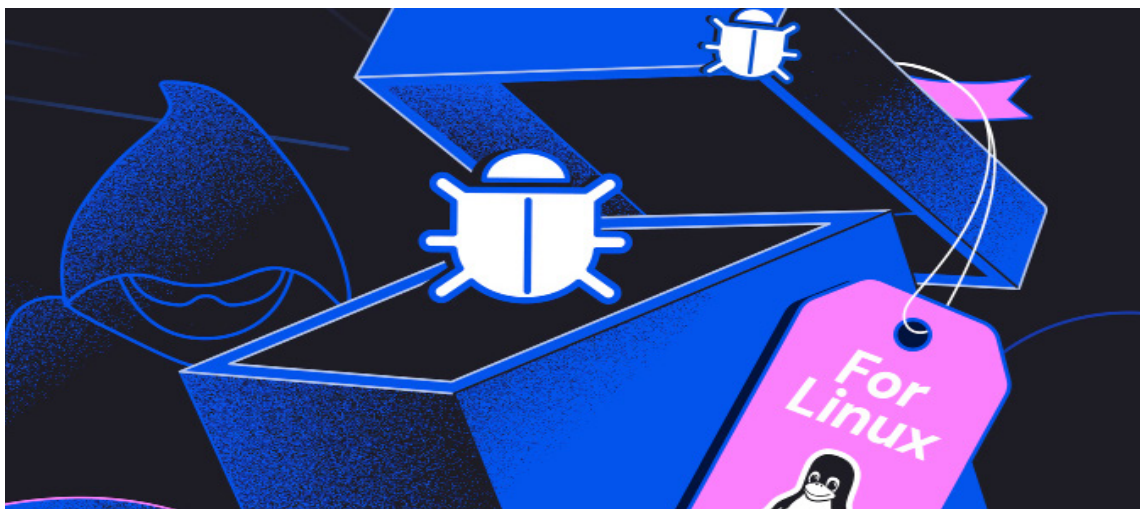

ULTRALYTICS AI LIBRARY Compromised for Cryptomining

A recent supply chain attack targeted the Ultralytics AI library, a popular open-source machine learning tool. The attackers managed to inject malicious code into the library's GitHub repository, covertly transforming it into a cryptomining tool. The malicious update, downloaded by unsuspecting developers, hijacked computing resources to mine cryptocurrency, slowing down affected systems and raising electricity costs for victims. Software supply chain attacks have become increasingly common, as attackers recognize the efficiency of compromising widely used development tools. By infiltrating open-source repositories, malicious actors gain access to a broad range of targets with minimal effort.

This particular attack on Ultralytics underscores the importance of verifying the integrity of software dependencies before integration.

To mitigate risks associated with software supply chain attacks, developers should implement strict code review processes, use dependency scanning tools, and verify software signatures before deployment.

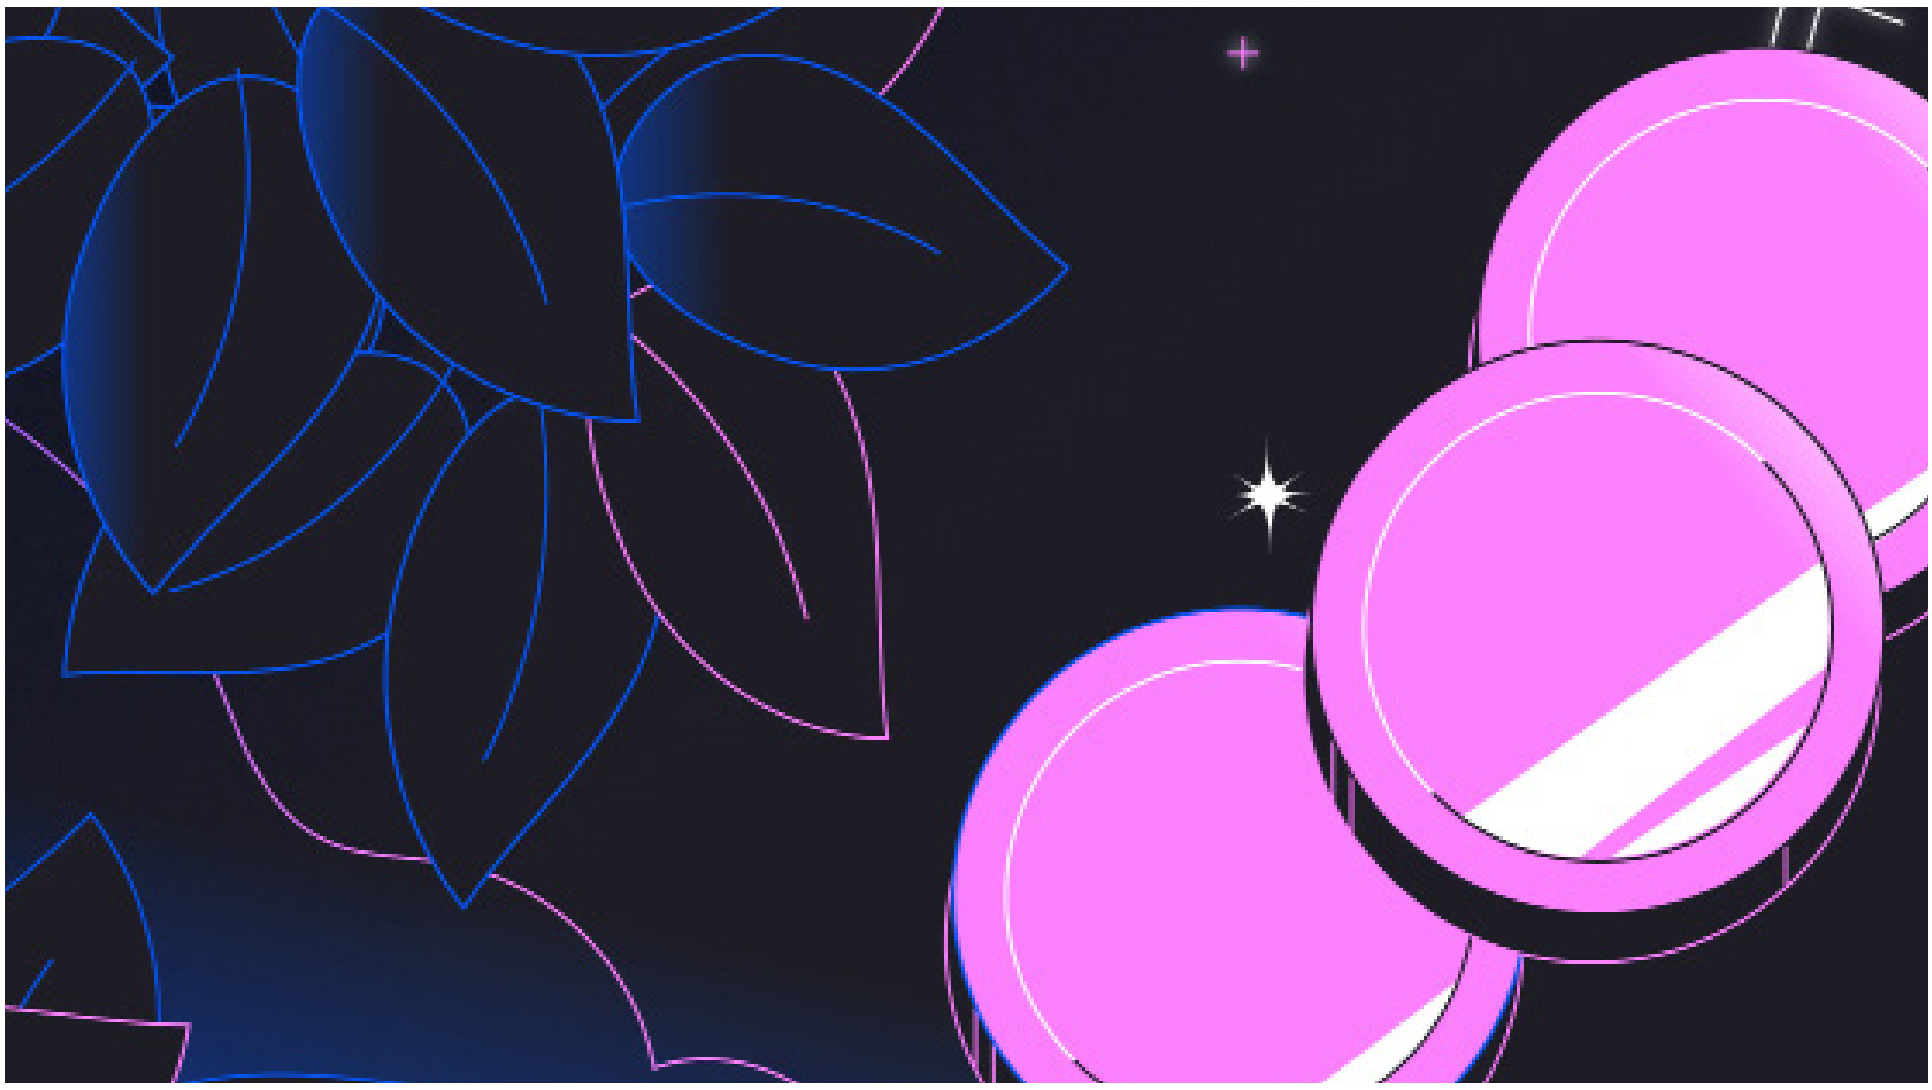

IVANTI ZERO – DAYS UNDER ACTIVE EXPLOITATION:

CVE-2025-0282 AND CVE-2025-0283

Two critical zero-days affecting Ivanti Connect Secure and Ivanti Policy Secure—CVE-2025-0282 and CVE-2025-0283—have been actively exploited in the wild. These vulnerabilities allow attackers to bypass authentication and execute arbitrary commands, targeting edge infrastructure widely deployed across enterprises.

Wiz Research identified active exploitation attempts, with attackers leveraging these flaws to drop webshells, establish persistence, and laterally move within networks. The vulnerabilities impact default installations and are especially dangerous given the appliance's location at the network perimeter. Ivanti released patches and mitigations, but threat actors had already gained footholds in some environments before fixes were issued.

Organizations should upgrade to patched versions immediately and inspect devices for signs of compromise. Apply additional controls such as segmentation and EDR visibility on edge devices to mitigate lateral movement.

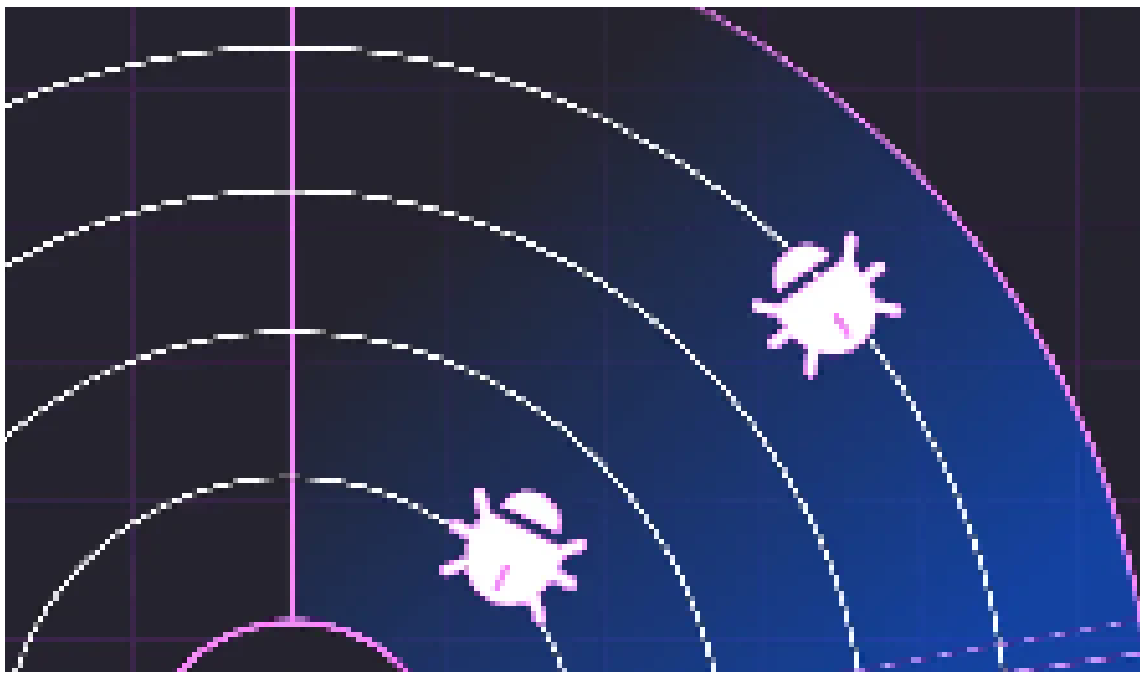

Nuclei Signature Bypass Allows UNDETECTED EXPLOITS

A vulnerability in the Nuclei scanner allowed attackers to bypass signature-based detections, leading to potential blind spots in automated security scans. The flaw stemmed from how YAML-based signature templates were validated and executed.

Wiz researchers demonstrated that by crafting malformed templates, attackers could inject payloads that appeared benign during scanning but executed malicious logic during runtime. This evasion technique could allow real vulnerabilities in web apps or APIs to remain undetected by security teams relying on Nuclei scans for coverage. Security teams using Nuclei should update to the latest release, validate custom templates, and combine static scanning with dynamic analysis to improve detection accuracy.

CVE-2024-0132:

Critical Vulnerability in NVIDIA AI Infrastructure

NVIDIA's AI inference server Triton was found to contain a critical vulnerability—CVE-2024-0132—that allows unauthenticated remote code execution. The issue lies in the server's handling of Python model files and could allow attackers to execute arbitrary commands on the host.

Exploitation requires uploading a malicious model archive, which the Triton server then deserializes using insecure logic. This could let attackers run code as root, pivot within a research environment, or exfiltrate sensitive AI datasets. Given the server's role in high-performance AI deployments, the risk to IP, models, and data integrity is substantial.

Organizations running NVIDIA Triton Inference Server should update to patched versions, limit upload permissions, and isolate inference workloads from broader infrastructure.

Spring Boot Actuator Misconfigurations Expose Credentials and Enable RCE

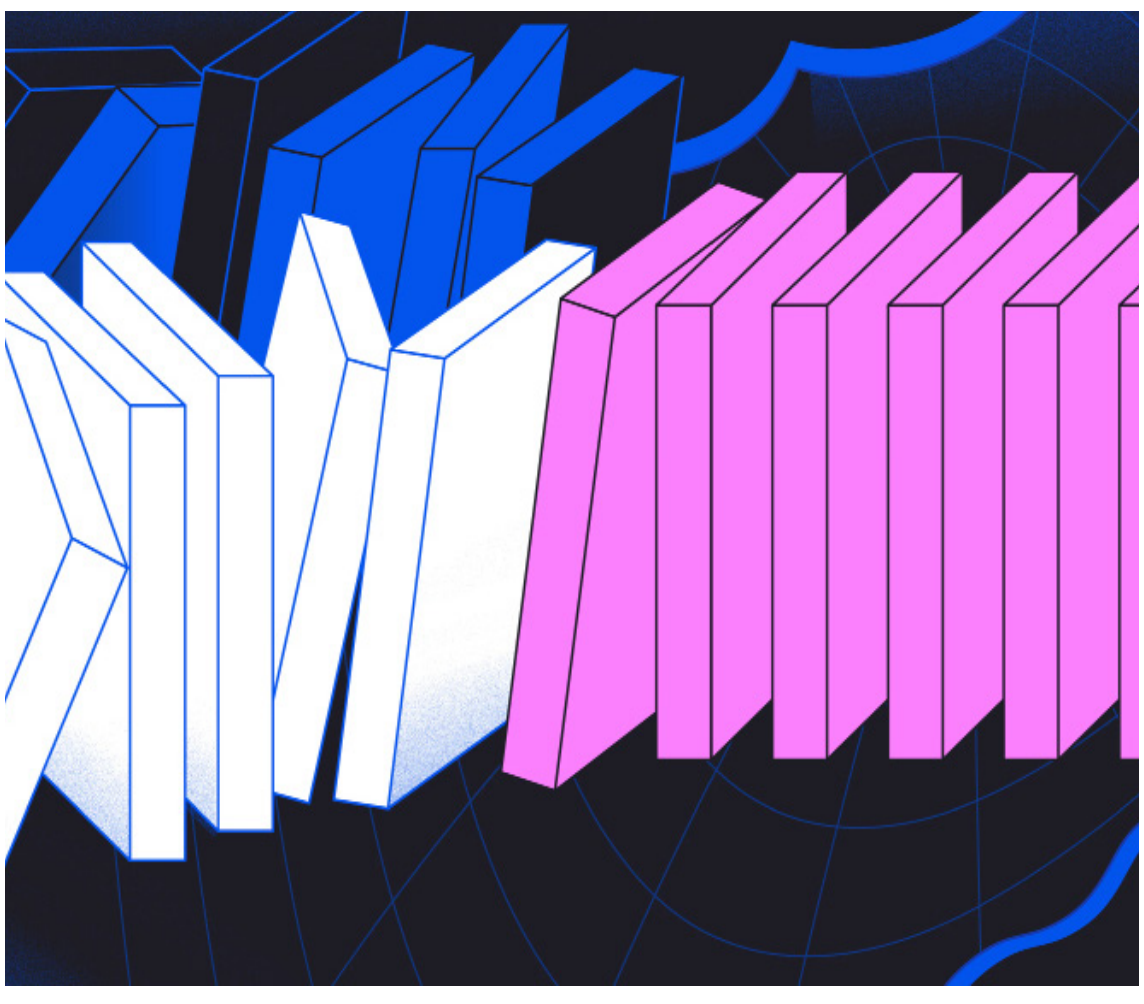

Wiz Threat Research uncovered misconfigurations in Spring Boot Actuator endpoints, affecting over 60% of analyzed cloud environments. These issues expose sensitive data such as environment variables, API keys, and credentials—and in some cases, allow remote code execution.

Critical risks include publicly accessible `/heapdump` endpoints containing plaintext AWS credentials, `/env` endpoints leaking configuration secrets, and `/gateway/routes` exposures that, when paired with vulnerable versions of Spring Cloud Gateway, could be chained into RCE (notably CVE-2022-22947). A notable real-world incident tied to such misconfiguration involved Volkswagen, where a heap dump inadvertently exposed access credentials.

Wiz found that 11% of cloud environments had exposed Actuator services, with nearly a quarter of those instances misconfigured. To reduce risk, organizations should disable non-essential endpoints in production, enforce authentication and access controls, and regularly audit for exposure.

"Over 60% of analyzed cloud environments had misconfigurations in Spring Boot Actuator endpoints, exposing sensitive data and enabling RCE."

PostgreSQL Exploited for Cryptomining in Cloud

Wiz Research uncovered cryptojacking campaigns abusing misconfigured PostgreSQL databases. Attackers leveraged the database's powerful extension and job scheduling features to deploy mining payloads.

The campaign exploited unsecured instances with superuser access, using `pg_execute_server_program` and custom extensions to install and persist mining software. Attackers also used obfuscation techniques to avoid detection by conventional EDR tools. Cloud teams should secure PostgreSQL with strong authentication, disable unnecessary superuser features, and monitor for abnormal CPU usage or suspicious SQL commands.

"Wiz Research discovered attackers exploiting misconfigured PostgreSQL servers to deploy cryptominers using superuser functions and obfuscation to evade detection."

PAN-OS CVE-2024-0012: Exploited in the Wild

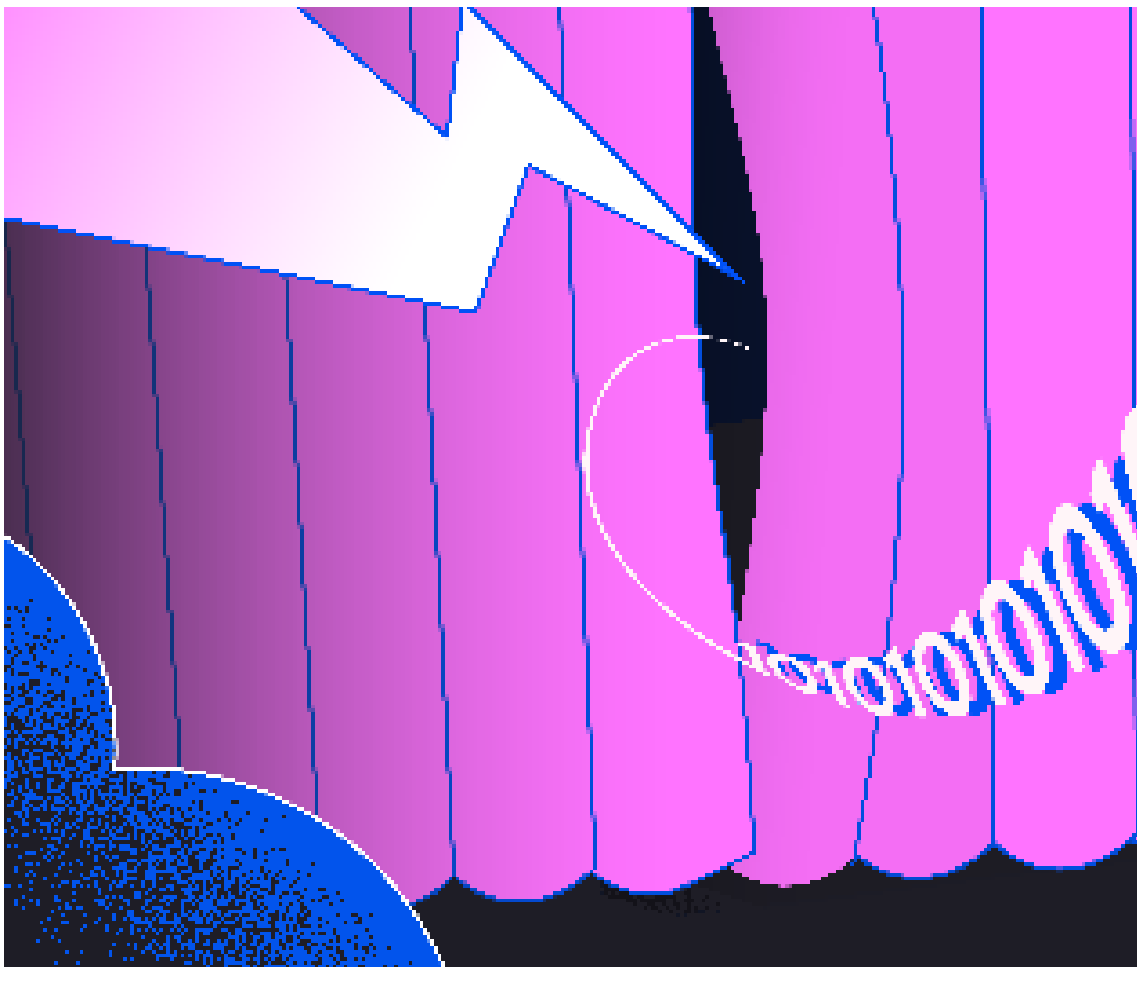

A severe security flaw in Palo Alto Networks' PAN-OS, CVE-2024-0012, is actively being exploited by cybercriminals. This zero-day vulnerability allows attackers to bypass authentication mechanisms and execute arbitrary commands on vulnerable firewalls. Given the critical role PAN-OS firewalls play in enterprise security, this exploit represents a significant threat to organizations.

Wiz Threat Research observed attackers using the vulnerability to manipulate firewall configurations, intercept network traffic, and create persistent access points. The high-profile nature of PAN-OS users, which includes government agencies, financial institutions, and cloud service providers, makes this vulnerability particularly alarming. Failure to address the issue promptly could lead to large-scale data breaches and network disruptions.

Security teams must act immediately by applying the latest patches and implementing additional security controls, such as network segmentation and continuous traffic monitoring.

Palo Alto Networks has released an emergency update, and organizations are strongly urged to apply it as soon as possible to prevent unauthorized access.

Ivanti EPMM Vulnerability Chain Enables Pre-Auth RCE (CVE-2025-4427 & CVE-2025-4428)

Wiz Threat Research has observed active exploitation of a critical vulnerability chain in Ivanti Endpoint Manager Mobile (EPMM), involving CVE-2025-4427 and CVE-2025-4428. When combined, these flaws allow unauthenticated remote code execution (RCE) on affected systems.

CVE-2025-4427 is an authentication bypass resulting from misconfigured route handling in Spring Security, exposing sensitive endpoints without proper access controls. CVE-2025-4428 is a post-authentication RCE vulnerability stemming from unsafe processing of user input in error messages, leading to Java Expression Language (EL) injection. Exploiting these vulnerabilities together enables attackers to execute arbitrary code without authentication.

Wiz researchers have identified multiple exploitation techniques in the wild, including the deployment of Sliver beacons, MySQL database dumps, web shell installations disguised as legitimate files, and reverse shell connections. These activities suggest involvement from various threat actors, some of whom have previously targeted similar infrastructure.

Organizations using Ivanti EPMM versions 11.12.0.4 and prior, 12.3.0.1 and prior, 12.4.0.1 and prior, and 12.5.0.0 and prior should urgently apply the patches provided by Ivanti. Additionally, implementing network-level restrictions on `/rs/api/v2/*` and `/mifs/rs/api/v2/*` endpoints is recommended until patches are applied.

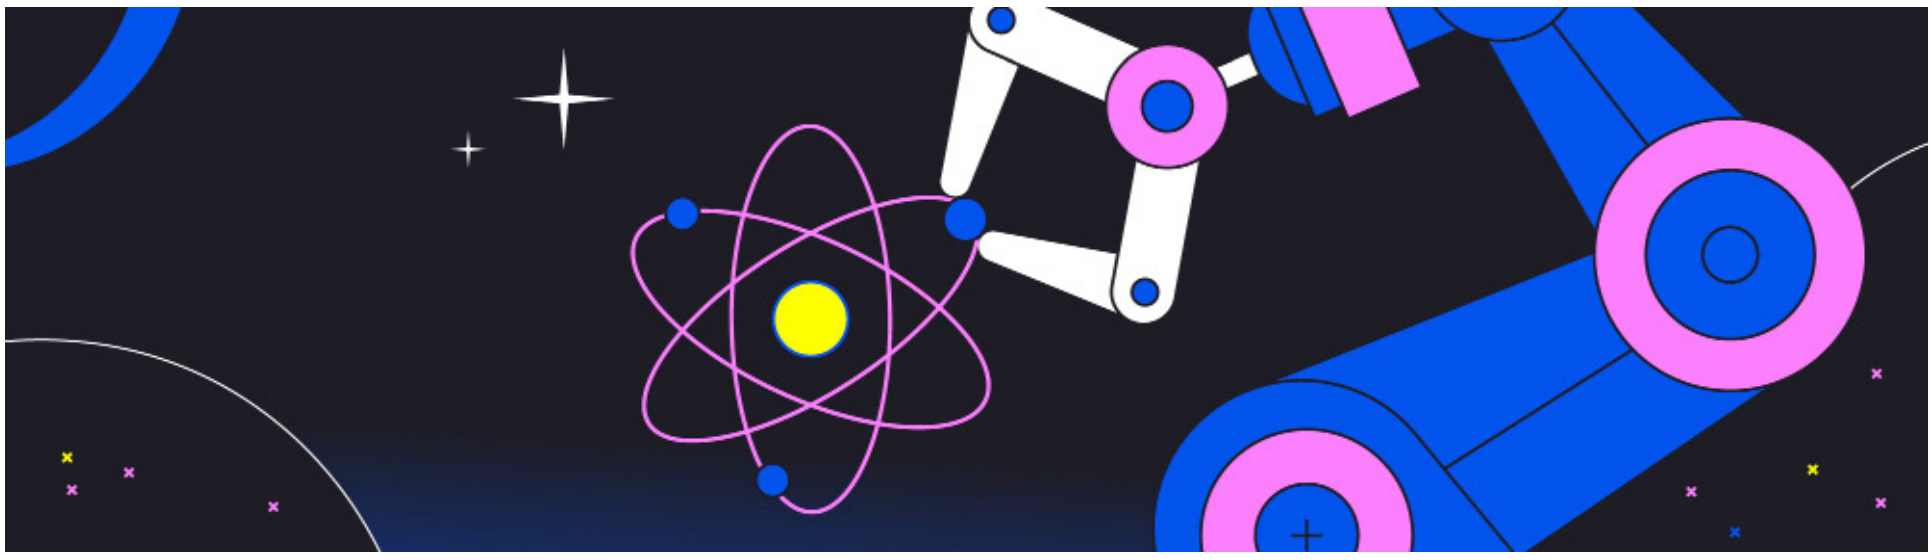

JINX-2401: Hijacking LLMs in AWS Environments

Wiz Research uncovered a new attacker dubbed JINX-2401, targeting permissive metadata access in AWS to hijack large language model (LLM) workloads. By compromising vulnerable AI workloads running on Amazon SageMaker or similar services, adversaries can extract models, manipulate outputs, or inject malicious prompts.

The attack chain leverages access to the instance metadata service (IMDS) and overly permissive IAM roles to pivot into other cloud services or exfiltrate sensitive data. Because LLMs often operate with elevated privileges and broad access to training data, this threat poses a serious risk to both data security and model integrity.

Organizations deploying LLMs in cloud environments should enforce strict IAM least-privilege policies, disable IMDS where unnecessary, and monitor for suspicious API activity tied to model endpoints.

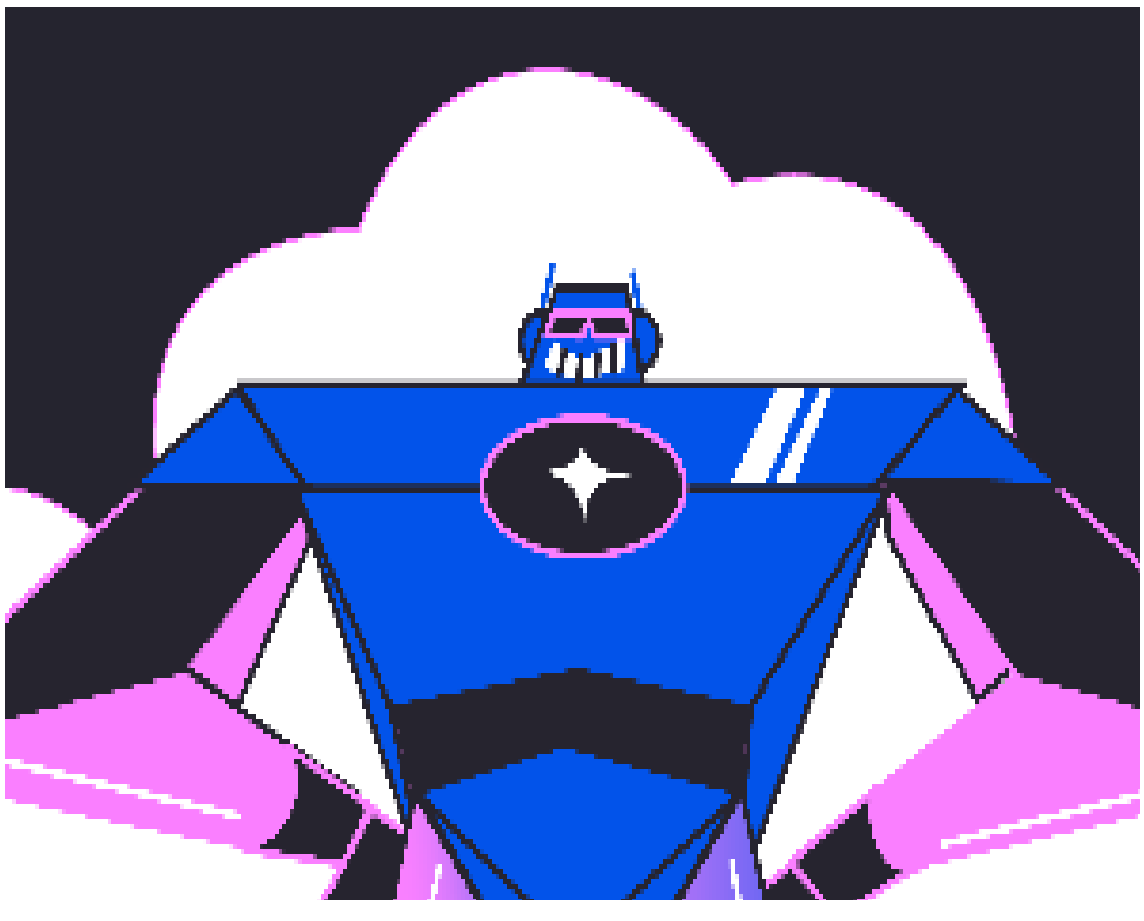

The Lighter Side of **Cloud Security**

Take a break. Solve a puzzle. Laugh a little.

R R Z R Q W Z C I E B P C Y L A P T K O O N Z
I Y C N E I U A P X Y M M A L I C I O U S C G
S G F S R K P L Y A D O R E Z A I Y F A Y O G
K A O U Y K S M L L D R H Z L M E P S S F P Z
G B N P T C Z R W D S M B R S U R U G N U T U
B A X P I M S B A E O S B O O J O X I W S P T
W W U L L U K Y Q Y C U F Y F J M X F Q J S J
I C M Y I Z O P Q K S S V C X H U N N D U L L
Z M A C B H C A G C G H D O E X P L O I T H B
K W T H A J Y S H A C S P M K R G G C G Y D S
B P P A R Y I S C T C X D K P E B K S H A N Z
F U E I E W B G A T E H P R P A P R I Y T N B
U Y G N N B G O E A J C V A O B T T M D R S O
K J R F L Q P E R X V M P X E C R C Z O Y E V
K T D Q U T Y B B T H V X T F U R H H M W U J
U J A A V O M Y C C Z J Z N Z V E K Z E V C M

Find the following words in the puzzle. Words are hidden vertically (up and down) horizontally (left and right), and diagonally (from top-left to bottom-right).

Patch
Vulnerability
CVSS
EPSS
POC
KEV
RCE
Bypass
Breach
Malicious
Attack
supply-chain
APT
CSPM
CVE
IAM
Risk
Wiz
Exploit
SOC
Zero-Day
Misconfig
LLM
K8s

CYBER *Horoscope*

Aries
Mar 21 – Apr 19

You're the fearless first responder- but slow down before enabling every new beta feature. Not everything shiny belongs in prod.

Taurus
Apr 20 – May 20

Your steady nature brings reliability to your org's posture. But don't resist updates too long—yes, even that one annoying patch

Gemini
May 21 – Jun 20

Two tabs open? Try 200. You thrive on multitasking, but beware of info leaks in your browser extensions. Curiosity is great— just sandbox it.

Cancer
Jun 21 – Jul 22

The protector of the digital home. You're emotionally invested in every alert—just don't let false positives ruin your weekend.

Leo
Jul 23 – Aug 22

Seeker with security swagger. You love sharing findings— just make sure your tweets don't violate any NDAs.

Virgo
Aug 23 – Sep 22

Detail-obsessed and vulnerability-averse. You saw that misconfigured bucket *before* the tool flagged it. Just don't get stuck fixing everyone else's code.

Libra
Sep 23 – Oct 22

The diplomat of incident response. You bring balance to the war room— but remember: not every stakeholder deserves equal say in a zero-day.

Scorpio
Oct 23 – Nov 21

Master of secrets (and secrets managers). You're intense, private, and three steps ahead. But don't forget to rotate those keys.

Sagittarius
Nov 22 – Dec 21

The explorer of unknown CVEs. You're always scanning for new adventures— just don't skip over the basics like MFA.

Capricorn
Dec 22 – Jan 19

Structured, strategic, and quietly running the tightest IAM policy around. Just remember: perfection is great— but "secure enough" ships faster.

Aquarius
Jan 20 – Feb 18

Your ideas are way ahead of their time— sometimes even your team's SIEM can't keep up. Keep innovating, but don't forget to document.

Pisces
Feb 19 – Mar 20

Intuitive and empathetic, you *feel* risk before it materializes. Just make sure your gut check is backed by data.

Researcher Recipe Feature

Merav's One-Bowl Chocolate Chip Cookies

Makes 12 big, chewy cookies

Craving something sweet? These classic chocolate chip cookies are quick to whip up—no mixer, no mess, just one bowl and a whole lot of comfort.

Ingredients:

- 150g soft butter
- ½ cup white sugar
- ½ cup brown sugar
- 1 large egg
- 1½ cups all-purpose flour
- ½ tsp baking soda
- ½ tsp baking powder
- 100g chopped milk chocolate
- 50g chopped dark chocolate

Instructions:

1. Mix the base: Soften the butter in the microwave (don't melt it completely!) and stir in both sugars until the mixture is creamy.

2. Add the egg: Make sure the butter mixture isn't too hot, then mix in the egg until smooth.

3. Dry ingredients: Gradually stir in the flour, baking soda, and baking powder. Mix until a soft dough forms.

4. Chocolate time: Fold in the chopped milk and dark chocolate until evenly distributed.

5. Chill: Cover the dough and refrigerate for at least 1 hour (up to 24). This step is key for that perfect chewy texture.

6. Shape & bake: Scoop the dough into 12 large balls and place them well-spaced on a lined baking sheet.

7. Into the oven: Bake at 180°C (350°F) for about 10 minutes, until the edges are just golden and the centers are still soft.

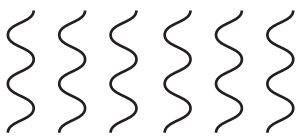

Let them cool slightly (if you can wait). Best enjoyed warm with cold milk—or straight off the tray, no judgment.

Scott's Buttermilk Pie

Ingredients:
- 3 eggs
- 1 1/4 cups sugar
- 2 tablespoons flour
- 1/2 cup melted butter
- 1 cup buttermilk
- 1 teaspoon vanilla extract
- 1 tablespoon lemon juice
- 1/2 teaspoon salt

Instructions:
Whisk all the ingredients together in a mixing bowl until smooth. Pour the mixture into a pie crust — I like using a store-bought graham cracker crust.

Bake at 350°F (175°C) for 45 minutes.

Best served warm (an hour or two after it comes out of the oven) with a scoop of vanilla ice cream. Pure heaven.

DATA MAZE

Secure the Cloud — If You Can!

Navigate the maze by choosing the right security moves. But beware: every wrong turn leads straight to a hacker. Only the sharpest minds will reach the safe cloud.