wallarm

API
THREATSTATS™
REPORT 2026

The New API Risk Multiplier

Contents

Executive Summary

If an attacker could control only one part of your infrastructure, they would pick your APIs. No one is surprised when you say that APIs are the backbone of modern digital business, but not everyone realizes that they've also become the most effective way in for attackers.

The findings are clear and consistent across every dataset we examined, no guesswork required. Attackers are not chasing exotic flaws. They are exploiting repeatable failures in identity, access control, and exposed interfaces, often at machine speed and massive scale. APIs dominate vulnerability disclosures, account for nearly half of confirmed exploited vulnerabilities, and sit at the center of the most consequential breaches of the year.

Three themes define API risk in 2025:

Abuse Beats Bugs

The API ThreatStats Top 10 shows attackers favor logic abuse, trust failures, and resource consumption over traditional code defects. Those behaviors surface most clearly as cross-site issues and broken access control, which rose to the top by attack volume. Injections remain a persistent, high-impact threat, not because attackers prefer them, but because APIs still process vast amounts of untrusted input at scale.

AI Security is API Security, Now with Consequences

36% percent of published AI vulnerabilities involve APIs. And 36% (yes, the same percentage) of AI-related exploited vulnerabilities also involve APIs. As AI adoption accelerates, APIs have become the primary interface attackers use to reach models, data pipelines, and automated workflows.

Autonomy Changes the Blast Radius

Model Context Protocol (MCP) emerged as a leading indicator of where API risk is heading. Although arguably still early in adoption, MCP accounted for 14% of published AI vulnerabilities, showed explosive growth during the year, and was involved in a Top 10 API breach. When APIs act on behalf of agents, failures no longer just expose data, they delegate authority.

The API ThreatStats report connects these trends end to end. Vulnerabilities that are fast, remote, and trivial to exploit predict how attackers operate. Exploited vulnerabilities in the CISA KEV catalog confirm APIs as the dominant surface, and that AI security really is API security. Breach analysis shows identity failures and exposed APIs, not advanced attackers, drive real-world impact.

For security leaders, the takeaway is direct, and actionable. Improving API security is not about chasing new attack classes or one-off fixes. It is about systematically addressing identity, exposure, and abuse before automation and scale turn familiar weaknesses into material business risk. The rest of this report explains how, and where, to focus next.

THREATSTATS SPOTLIGHT

MCP

Model Context Protocol: When APIs Become Autonomous

Model Context Protocol (MCP) is still new enough that many teams think of it as an architectural detail, not a real attack surface, but the 2025 data shows that assumption is already out of date. MCP has rapidly emerged as a concentrated source of API risk across vulnerabilities, exploits, and breaches, especially in agentic AI environments where APIs do more than just respond to requests. They act.

AI VULNERABILITIES IN 2025

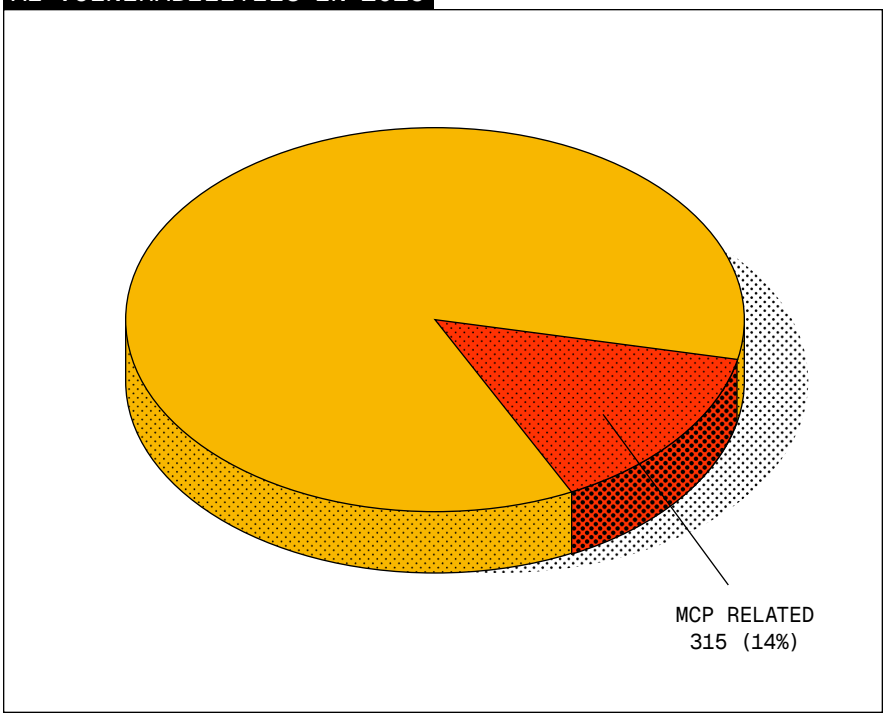

MCP represents a rapidly emerging API attack surface with outsized vulnerability growth despite early adoption

MCP IN VULNERABILITY DATA

A Small Surface with Outsized Risk

In 2025, we identified 315 MCP-related vulnerabilities, representing 14% of all AI vulnerabilities published during the year. Normally, we would want to see a year-over-year comparison, but in the 2024 data, MCP wasn't prevalent enough to measure separately. We did look at MCP specifically in our Q3 report, and calculated backwards to Q2 to determine a growth trajectory. Our Q3 API ThreatStats report showed a 270% increase in MCP vulnerabilities from Q2 to Q3. This is a striking concentration and stunning momentum for a protocol that is still early in adoption.

MCP is effectively a control plane API for agents, and when that plane is exposed or misconfigured, attackers gain leverage over autonomous workflows rather than single endpoints.

MCP vulnerabilities are critical because they consistently combine three failure modes:

1. OVER-PERMISSIONED TOOLS, WHERE AGENTS ARE GRANTED BROAD API ACCESS BY DEFAULT.

+

2. DIRECT API EXPOSURE, OFTEN WITH TYPICAL API VULNERABILITIES LIKE WEAK AUTHENTICATION OR OVERLY BROAD SCOPE.

+

3. LACK OF RUNTIME ENFORCEMENT, MEANING POLICY VIOLATIONS ARE ONLY VISIBLE AFTER DAMAGE OCCURS.

While MCP vulnerabilities are often bucketed into "AI issues," the underlying mechanics are firmly API-centric. MCP is effectively a control plane API for agents, and when that plane is exposed or misconfigured, attackers gain leverage over autonomous workflows rather than single endpoints.

MCP IN EXPLOITATION

APIs as the Path to Agent Control

The exploit data reinforces the vulnerability pattern. In the CISA KEV analysis in this report, 36% of AI-related exploited vulnerabilities also exposed an API attack surface. MCP-related flaws fit squarely into this overlap.

When MCP vulnerabilities are exploited, attackers are not chaining exotic AI-specific techniques. They are abusing familiar API mechanics, like unauthenticated access, path traversal, unsafe backend calls, to reach agent tooling and execution environments. The difference is impact: MCP turns API compromise into autonomous execution, allowing a single exploit to trigger cascading actions without further attacker interaction.

This is how MCP meaningfully changes the risk calculus. APIs have always been attractive because they are fast and remote. MCP adds persistence and agency to that equation.

MCP IN BREACHES

From Theory to Incident

MCP risk is no longer hypothetical. One of the Top 10 API breaches of 2025, detailed later in this report, involved thousands of exposed MCP servers, where an API-reachable path traversal flaw allowed attackers to access agent infrastructure used in production AI workflows.

This incident followed the same pattern seen elsewhere in the report:

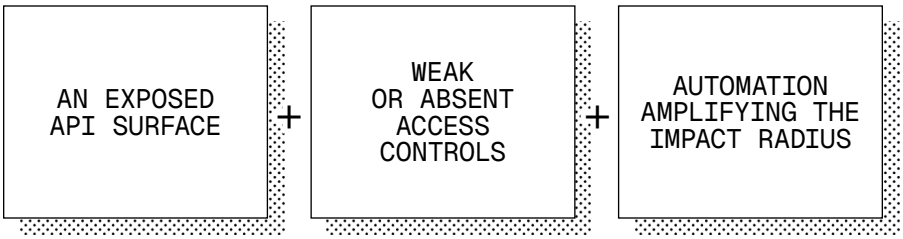

What makes MCP distinct is that the compromised APIs were not simply serving data. They were orchestrating tools and actions on behalf of agents, dramatically increasing the potential impact of the API exposure.

Why MCP Deserves a Spotlight

Across the vulnerabilities, exploits, and breaches we analyzed for this report, MCP shows a consistent signal:

1. IT REPRESENTS A DISPROPORTIONATE SHARE OF AI-RELATED API VULNERABILITIES

+

2. IT CAN BE EXPLOITED THROUGH STANDARD API ABUSE TECHNIQUES, NOT NOVEL AI ATTACKS

+

3. WHEN BREACHED, IT ENABLES AUTONOMOUS, CASCADING IMPACT RATHER THAN ONE-OFF ACCESS

MCP is an early indicator of where API risk is heading. As APIs increasingly act on behalf of systems instead of users, failures in authentication, authorization, and runtime control no longer just expose data. They delegate power to attackers.

Simply put, securing agentic AI starts with securing the APIs that give agents their context, tools, and authority. And MCP is the primary component in this next-generation infrastructure.

Simply put, securing agentic AI starts with securing the APIs that give agents their context, tools, and authority.

API THREATSTATS TOP 10

How API Risk Evolved Across 2025

The Wallarm API ThreatStats Top 10 shows how attackers actually target APIs in production, based on what we see happening, not what looks risky on paper. Unlike lists based on disclosure volume or theoretical risk, this ranking is derived from observed attack activity and shows where adversaries consistently focus their effort.

This section is anchored on quarter-over-quarter attack volume data across 2025. Each quarter reflects the relative intensity of abuse across API weakness categories, revealing which risks are persistent, which accelerate quickly, and which emerge as architectures change.

TOP 10 API RISKS (YOY CHANGE)

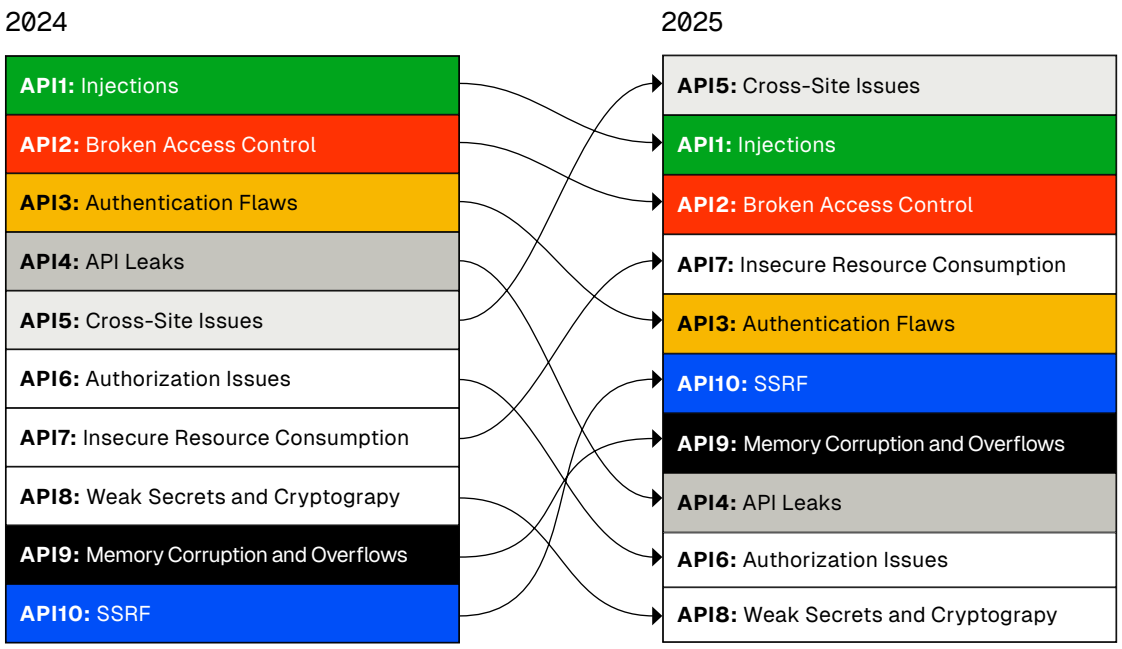

TOP 10 API RISKS (QOQ CHANGE)

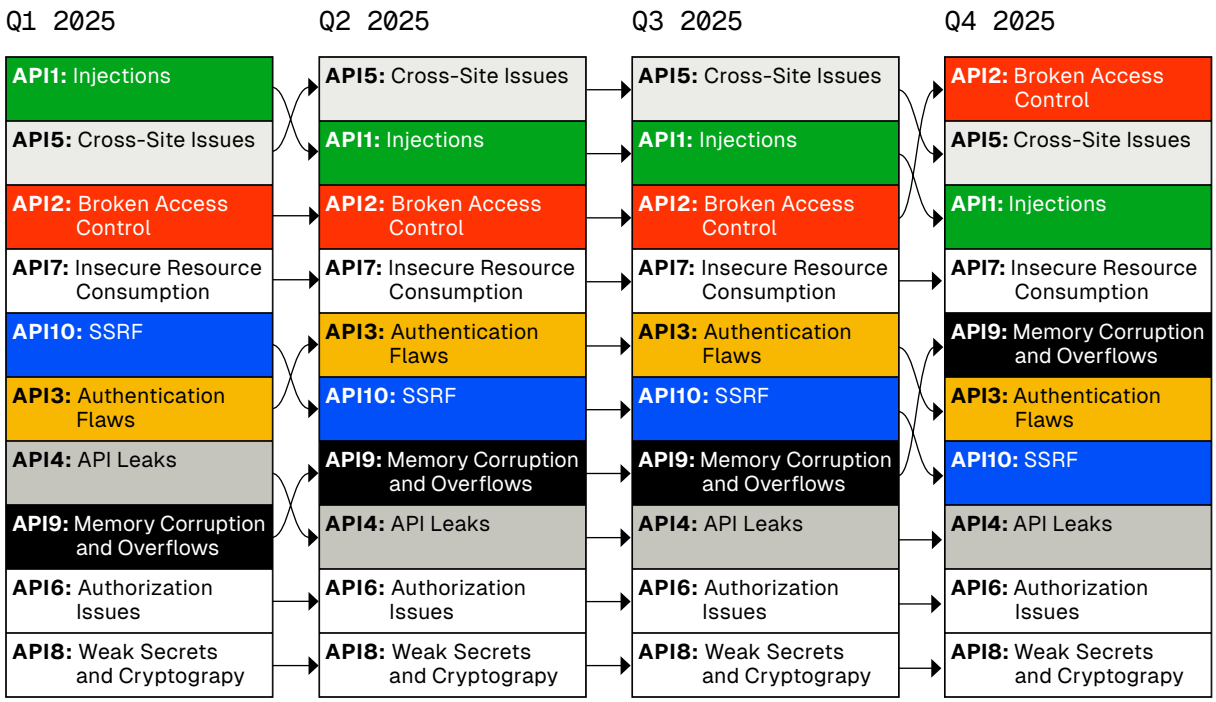

How to Interpret the ThreatStats Top 10

Each category represents a class of API weakness defined by attacker behavior, not code taxonomy. Rankings reflect observed attack volume, meaning movement up or down the list signals real changes in attacker focus rather than shifts in disclosure practices.

Categories that remain near the top across all quarters indicate structural weaknesses in how APIs are designed and operated. Categories that shift rapidly highlight where attackers are adapting tooling, automation, or monetization strategies.

THE 2025 LEADERBOARD

What the Data Shows

Cross-Site Issues emerge as the dominant API threat in 2025. While they ranked mid-pack in 2024, they moved decisively upward and finished the year as the most abused category overall.

This rise reflects how effectively cross-site techniques exploit modern API-driven user experiences. These attacks are easy to automate, require little reconnaissance, and weaponize implicit trust between browsers, sessions, and APIs, making them especially effective against consumer-facing platforms.

Injections Remain a Constant Anchor Threat

While Injections took the top spot in 2024, they ranked second overall in 2025 and never fell below the second position in any quarter. It's clear that despite years of industry education about injections, APIs continue to process vast volumes of untrusted input and pass it directly into downstream systems.

It's clear that despite years of industry education about injections, APIs continue to process vast volumes of untrusted input and pass it directly into downstream systems.

The expansion of AI-driven APIs amplifies the impact of injection flaws. A single successful injection can now influence models, data pipelines, or automated decisions at machine speed. The data confirms the uncomfortable truth that injection vulnerabilities should not be treated as legacy attacks; they are supremely relevant.

Broken Access Control Continues to Enable Scale

Broken Access Control holds the third position for the year, down from the number two spot last year. It was consistently third with a bump to first in Q4.

When attackers discover new enumeration or privilege-escalation paths, exploitation scales rapidly. Access control remains one of the most fragile controls in distributed API environments.

THE MIDDLE OF THE CHART

Where Attackers Optimize

Insecure Resource Consumption Climbs Steadily

Insecure Resource Consumption finished fourth overall in 2025, up from seventh in 2024. Notably, it held the same relative position across all four quarters, signaling consistent attacker interest rather than short-lived spikes.

Scraping, enumeration, and denial-of-service-style abuse continue to grow as attacker tooling becomes more automated and increasingly AI-assisted. APIs without effective rate limits, quotas, and behavioral controls remain easy targets.

Authentication Flaws and SSRF Trade Places

Authentication Flaws and Server-Side Request Forgery occupy the middle of the Top 10, reflecting persistent weaknesses rather than volatility. Weak token handling, long-lived credentials, and insufficient identity enforcement keep authentication failures highly exploitable, while SSRF continues to expose backend connectivity and cloud metadata paths. It's notable that when compared to 2024, Authentication Flaws fell from spot number three to five and SSRF rose from tenth to sixth.

CHRONIC BUT DANGEROUS

Always on the Board

Several categories remain lower in total volume but are consistently present.

Memory Corruption and Overflows

Memory Corruption and Overflows, introduced in last year's report, climbed from ninth to seventh, signaling a meaningful shift toward binary and high-performance APIs, driven by AI and GPU-accelerated workloads.

API Leaks

API Leaks fell from fourth to eighth, persisting due to exposed endpoints, forgotten parameters, and shadow APIs that evade inventory.

Authorization Issues

Authorization Issues took the ninth spot, lower in volume, but reflecting chronic difficulty in managing roles and permissions at scale.

The continued presence of these issues reinforces that API risk is cumulative. These weaknesses rarely appear alone and often compound during real attacks.

What the Quarter-over-Quarter Progression Reveals

Taken together, the 2025 ThreatStats Top 10 tells a pretty clear story:

1. ATTACKERS NOW FAVOR ABUSE OVER BUGS, TARGETING LOGIC, TRUST, AND USAGE PATTERNS.
2. AI AMPLIFIES EXISTING WEAKNESSES RATHER THAN INTRODUCING ENTIRELY NEW ONES. MORE APIS, MORE OPPORTUNITIES FOR ABUSE.
3. RUNTIME BEHAVIOR DEFINES API RISK, NOT PRE-PRODUCTION TESTING ALONE.

The API ThreatStats Top 10 exists to close the gap between theoretical risk and operational reality. It shows not just what could go wrong, but what attackers are already exploiting, quarter after quarter. It's a useful lens through which to review the remainder of the API ThreatStats report.

API Vulnerability Trends

We analyzed as close to a complete set of published vulnerabilities from 2025 as possible, 67,058 security bulletins in total. If you're a repeat reader of the API ThreatStats report, this year's findings reinforce a now-familiar theme (perhaps with sharper edges): APIs remain the dominant vulnerability surface, and AI has accelerated both the volume and impact of API flaws.

From the data set, in total, 11,053 vulnerabilities (17%) were API-related, and 2,185 vulnerabilities (3%) were AI-related. Critically, 786 vulnerabilities overlapped both categories, meaning more than one-third (36%) of all AI vulnerabilities are also API vulnerabilities. In practical terms: if you cannot secure your APIs, you cannot secure your AI. Our 2025 report was titled AI Security is API Security, and the data shows that statement is no longer provocative; it's just descriptive.

API VULNERABILITIES

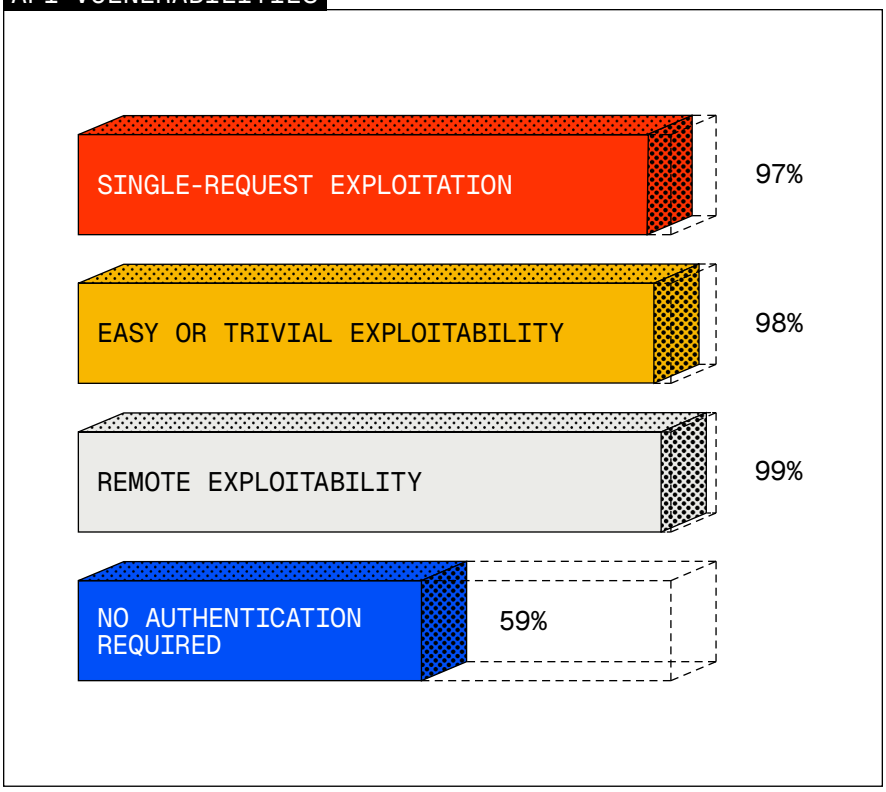

Most API vulnerabilities are fast, remote, and trivial to exploit, leaving little room for delayed detection

The 2025 data shows that API vulnerabilities overwhelmingly favor speed, simplicity, and scale from an attacker's perspective. More importantly, these characteristics explain why traditional security controls consistently fail against API abuse.

Single-request attacks dominate the vulnerability landscape.
Fully 97% of API vulnerabilities can be exploited with a single request. Multi-step or chained attacks are statistically insignificant in the vulnerability data. This matters because most security programs are still optimized for patterns over time: session analysis, behavioral baselines, or correlation across multiple requests. When the attack completes in a single call, defenders don't get a second chance. Prevention has to happen inline and per request, or it doesn't happen at all.

Exploitability is trivial, which removes skill as a limiting factor.
When exploitation is this easy, scale, not sophistication, becomes the dominant risk factor. An overwhelming 98% of published API vulnerabilities are rated easy or trivial to exploit. Only about 1% require advanced skill or complex conditions. This eliminates attacker sophistication as a meaningful constraint. When exploitation is this easy, scale becomes the dominant risk factor. Attackers don't need to be clever; they just need automation.

When exploitation is this easy, scale, not sophistication, becomes the dominant risk factor.

Attackers don't need to be clever; they just need automation.

Remote exploitation is the default operating condition.
Nearly 99% of API vulnerabilities are remotely exploitable, which shouldn't be a surprise because APIs are inherently designed for remote connections. It's worth keeping this fact in mind, however. Physical proximity, local access, or internal footholds are largely irrelevant for API vulnerabilities. From a risk perspective, this means that internet exposure equals attack exposure.

Authentication gaps turn APIs into anonymous attack surfaces.
In 59% of API vulnerabilities, no authentication is required at all. When endpoints are callable without identity, every other control becomes harder: rate limiting, authorization, abuse detection, and attribution. Anonymous access dramatically lowers attacker cost while increasing defender uncertainty.

In 59% of API vulnerabilities, no authentication is required at all.

Public exploit code accelerates weaponization.
30% of published API vulnerabilities have public exploit code available. This shortens the window between disclosure and active exploitation to near zero. But exploit code also aids in detection. When 70% of these published vulnerabilities don't have exploit code, it means that tool vendors have to work that much harder to build reliable, accurate detection. Simple signatures just won't cut it.

Business Impact and Threat Actors

These exploit characteristics may be easy to gloss over, but understanding them is critical to building adequate defenses. Organizations require tools that can stop highly automated exploits at scale. After-the-fact analysis and detection is too slow to be effective.

IMPACT VS ATTACKER SKILL

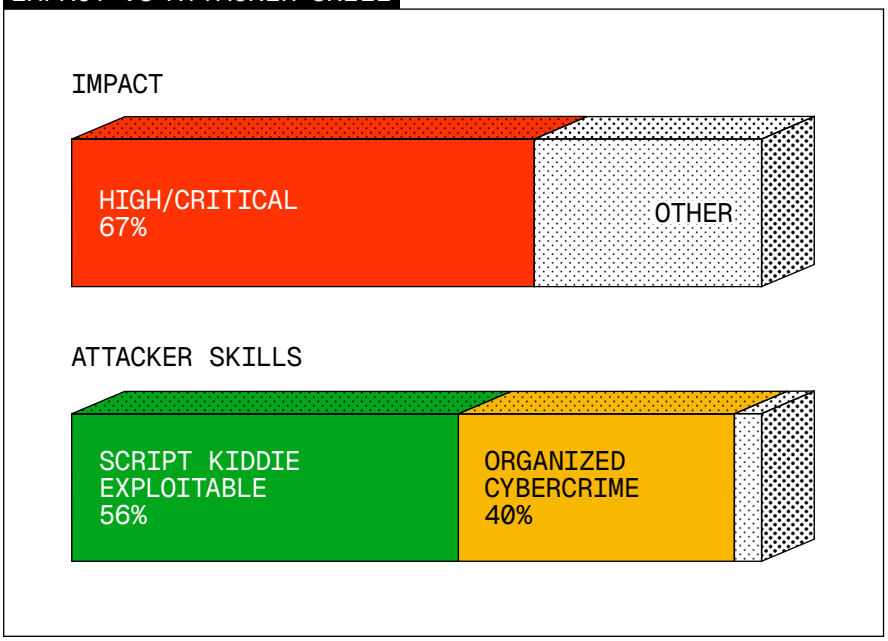

Most high-impact API vulnerabilities do not require advanced attackers

67% of API vulnerabilities are rated High or Critical. More than 56% are exploitable by script kiddies, with an additional 40% associated with organized cybercrime. Nation-state activity exists, but it is statistically marginal.

This combination of high impact, low skill, and massive scale is what makes API risk uniquely persistent.

This combination of high impact, low skill, and massive scale is what makes API risk uniquely persistent.

OWASP API Top 10 (2023) Distribution

The distribution of vulnerabilities across the OWASP API Top 10 demonstrates that risk concentrates around a small number of systemic control failures, not a long tail of edge cases.

Broken Object Level Authorization (API1) leads the dataset at 39%, followed closely by Unrestricted Resource Consumption (API4) at 35% and Server-Side Request Forgery (API7) at 33%. These categories frequently co-occur, reflecting how real-world API failures cluster rather than appear in isolation.

When compared to the API ThreatStats Top 10, this OWASP distribution helps explain why certain attack categories dominate in production. OWASP highlights where control failures exist at design and implementation time, while ThreatStats shows how attackers are operationalizing those failures at runtime. For example, widespread object-level authorization gaps and unsafe resource handling map directly to the rise of broken access control and insecure resource consumption in the ThreatStats rankings. In other words, OWASP shows the structural weaknesses, and ThreatStats shows how attackers exploit them at scale.

LEADING API VULNERABILITY CATEGORIES

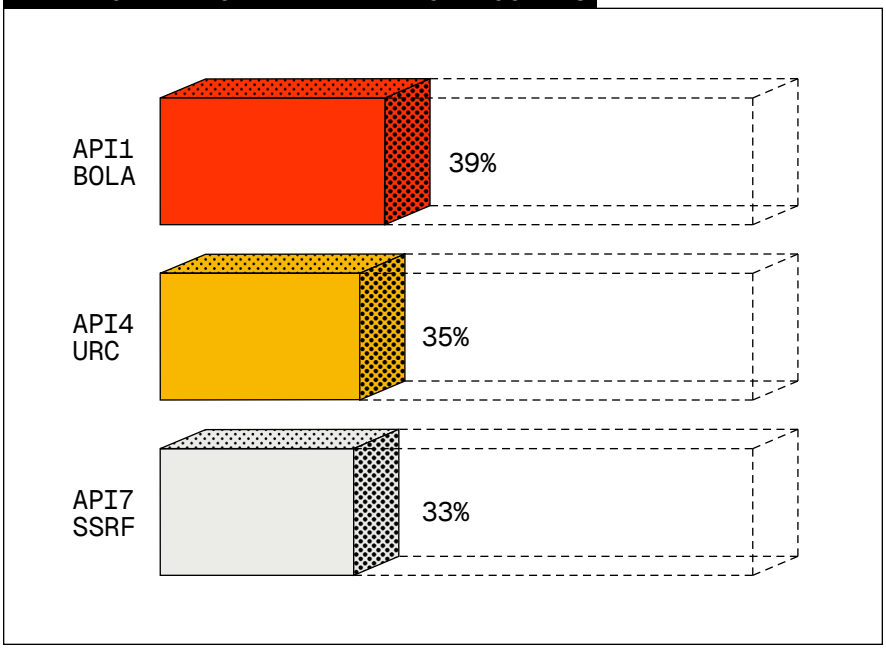

Real-world API vulnerabilities cluster around a small number of systemic failures

API1: Broken Object Level Authorization
API4: Unrestricted Resource Consumption
API7: Server-Side Request Forgery

API Protocol and Architecture Trends

REST APIs account for 96% of vulnerabilities, reinforcing (unsurprisingly) that this mainstream protocol drives the majority of risk. Authentication failures cluster around token-based schemes (JWTs, bearer tokens) rather than legacy methods, providing some hope that newer authentication methods dominate, though they are still vulnerable.

The API, AI Vulnerability Convergence

These reports are always more interesting when they identify trends. In the 2025 Annual API ThreatStats Report, Wallarm identified 439 AI vulnerabilities in 2024, and found that 77% were directly API-related. The 2026 analysis shows only 36% of AI vulnerabilities directly overlap with API vulnerabilities. While the overlap percentage decreased, the volume of AI vulnerabilities exploded. AI vulnerabilities grew from 439 to 2,185, a 398% increase year over year. At the same time, the raw number of vulnerabilities that are both AI and API-related increased from 439 to 786, a 79% increase. In short, even as overlap decreased, the real-world API attack surface tied to AI grew dramatically, and when API and AI weaknesses overlap, the same exploit characteristics dominate: remote access, single-request attacks, and consistently high business impact.

AI vulnerabilities grew from 439 to 2,185, a 398% increase year over year.

API/AI VULNERABILITY OVERLAP

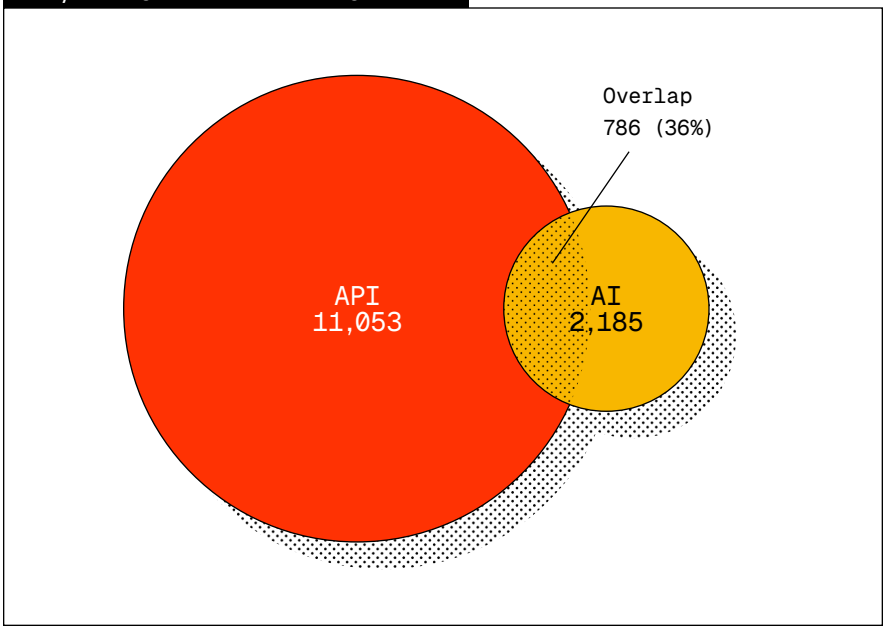

More than one-third of AI vulnerabilities are also API vulnerabilities

AI VULNERABILITIES

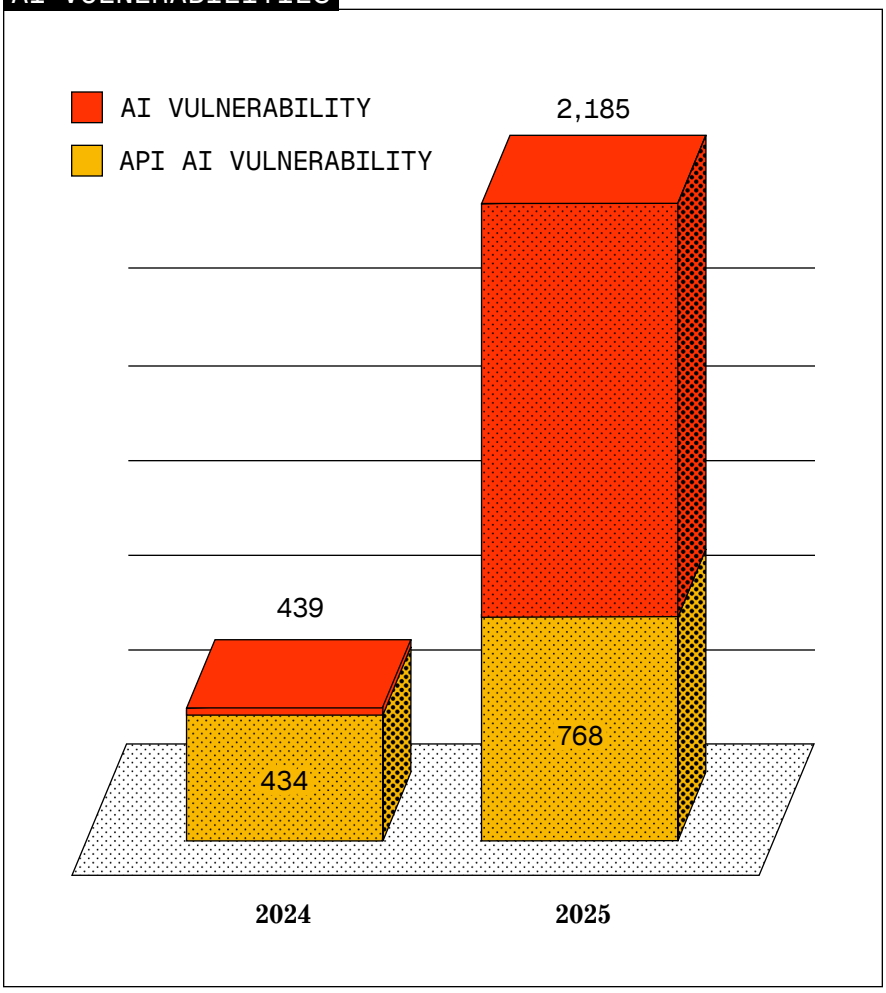

Even as definitions tightened, the number of AI vulnerabilities, and those rooted in API exposure, increased sharply year over year

EMERGING RISK

Model Context Protocol (MCP)

We've covered MCP in depth in the ThreatStats Spotlight, but part of that analysis is rooted in the vulnerability data here. 315 MCP-related vulnerabilities were identified in 2025, representing 14% of all AI vulnerabilities. MCP vulnerabilities frequently combine tool over-permissioning, insecure API exposure, and lack of runtime enforcement, making them especially dangerous in agentic AI deployments.

Vulnerability Trends Summary

The 2025 vulnerability data shows that APIs are the most consistently affected surface in published disclosures, largely because they are ubiquitous, externally reachable, and tightly coupled to business logic.

AI does not introduce an entirely new class of vulnerabilities in this dataset. Instead, it increases the frequency and impact of familiar API failure modes by expanding the number of exposed interfaces, the degree of automation, and the consequences of a single exploitable flaw.

Taken together, these trends indicate that API-related weaknesses dominate both traditional and AI-driven vulnerability disclosures, setting the conditions that make exploitation and downstream breaches more likely in later stages of the attack lifecycle.

API Exploit Trends

While the API ThreatStats Top 10 is built from attacks that Wallarm has actually observed, CISA's Known Exploited Vulnerabilities (KEV) catalog is one of the cleanest external signals we have for "this isn't theoretical, attackers are actively using it." For this section of the report, we analyzed KEV entries added during calendar year 2025 (n=245) and enriched them with NVD fields, including vulnerability descriptions, CWE mappings, CVSS context, and reference data.

Each CVE was then categorized using a deterministic methodology combining keywords, product context, CWE indicators, and CVSS network-exposure signals to determine whether the exploited vulnerability was API-related, AI-related, both, or neither.

APIs Dominate Exploited Vulnerability Activity

Based on this classification, 106 of the 245 KEVs added in 2025 (43%) are API-related. This makes APIs the single most common exploited surface in the KEV dataset.

KEV DISTRIBUTION

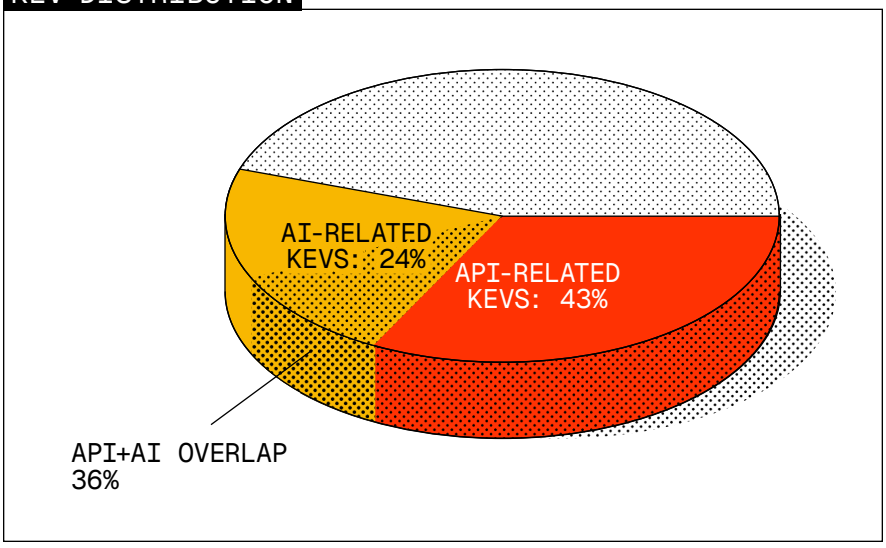

APIs represent the largest share of exploited vulnerabilities in 2025

AI-related exploited vulnerabilities are also material, accounting for 58 of 245 entries (24%). More importantly, 21 of those 58 AI-related KEVs (36%) are also API-related. This mirrors the vulnerability trends data, where 36% of all published AI vulnerabilities in 2025 were also API vulnerabilities.

The identical overlap percentages in vulnerabilities (36%) and exploited vulnerabilities (36%) are not coincidental. They indicate that the API, AI convergence observed at disclosure time persists through real-world exploitation. In other words, AI-related weaknesses do not become less API-centric as attacks progress. They become more operationally visible through APIs, which serve as the primary interface attackers abuse at scale.

Even when the root issue is framed as "AI," the exploit path frequently runs through APIs: model-serving endpoints, orchestration layers, integrations, and management interfaces. In practice, attackers are not inventing new techniques for AI systems; they are reusing the highest-volume API abuse patterns already proven at scale and applying them to AI-driven architectures.

Clearly, exploited AI risk is very often exploited API risk.

API Exploits Skew Remote and Low-Friction

API-related KEVs show stronger internet exploitability signals than the overall KEV population:

NETWORK ATTACK VECTOR

82 /106 (77%)

vs. 51% across all 2025 KEVs

NO PRIVILEGES REQUIRED

62 /106 (58%)

vs. 44% across all 2025 KEVs

When compared to the vulnerability analysis earlier in this report, these exploitation characteristics are not surprising. The vast majority of API vulnerabilities were already remotely exploitable, easy to abuse, and frequently exposed without authentication. The KEV data confirms that attackers are not selectively exploiting rare or complex edge cases. Instead, they are reliably exploiting the same low-friction conditions that dominate the vulnerability dataset. In practice, the properties that make API vulnerabilities common are the same properties that make them consistently exploitable in the wild. When exploitation touches APIs, it disproportionately resembles the attacker's ideal conditions: remotely reachable, low-to-no authentication, and immediately usable.

AI EXPLOIT TRENDS

Similar to APIs

AI-related KEVs skew remote as well, though with slightly more ambiguity in scoring inputs:

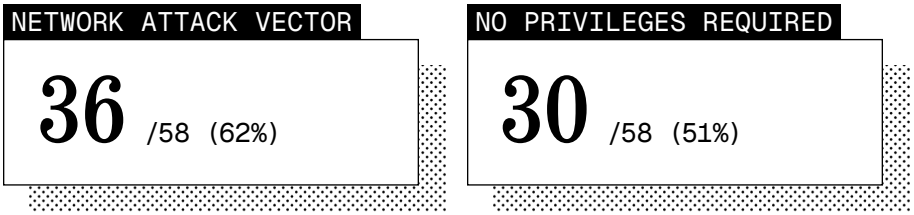

This reflects how AI systems are often deployed behind platforms, gateways, or internal networks until exposure creeps outward through APIs, integrations, and management layers. When that happens, exploitation accelerates quickly.

Overlap Between API and AI Exploitation

More than one-third of AI-related exploited vulnerabilities in 2025 also exposed an API attack surface. This overlap is not accidental. As noted elsewhere in this report, APIs are the control plane for AI systems: they expose models, tools, data pipelines, and automation hooks. AI risk materializes immediately when those API surfaces are vulnerable.

Top Vendors for API-Related Exploits

Among API-related KEVs in 2025, exploited vulnerabilities clustered around a relatively small group of platform and infrastructure vendors:

TOP INFRASTRUCTURE VENDORS

Microsoft	10 (9%)
FORTINET	7 (7%)
cisco	7 (7%)
ivanti	6 (6%)
ORACLE	5 (5%)

The next tier (ASUS, SonicWall, Gladinet, TP-Link, D-Link) each accounted for 3% of API-related KEVs.

This concentration reflects where commercially published APIs are most exposed and most operationally valuable: management planes, infrastructure services, and platforms deeply embedded in enterprise environments. It's important to keep the limitations of the KEV data set in mind here. The data here doesn't cover APIs that are made available as services, where a published security bulletin and mass exploitation is less likely to exist.

What the Exploit Data Tells Us

Three patterns stand out clearly from the 2025 KEV analysis:

1. APIS ARE THE DOMINANT EXPLOITED SURFACE. NEARLY HALF OF ALL NEWLY ADDED KEVS INVOLVE APIS.
2. EXPLOITED API VULNERABILITIES FAVOR SPEED AND REACH. REMOTE ACCESS AND LOW PRIVILEGE REQUIREMENTS ARE THE NORM.
3. AI EXPLOITATION DEPENDS ON APIS. WHEN AI SYSTEMS ARE COMPROMISED, APIS ARE THE DOORWAY.

For defenders, these conclusions reinforce an important truth; if your API security controls only work in theory, or only after the first request, they are already too late.

API Breach Trends

So far, we've covered vulnerability and exploit trends for 2025. The third leg of the ThreatStats stool is where those trends land in the real world: breaches. API breaches in 2025 followed a familiar script with a few new plot twists. The classics still dominated (broken authentication, leaked secrets, and exposed endpoints), but the supporting cast changed.

AI platforms and agentic ecosystems (MCP servers, LLM toolchains, API-driven automation) appeared frequently, and attackers treated them like any other high-value API surface: steal a token, find an admin endpoint, automate at scale.

This section analyzes 60 API-related breach incidents disclosed in 2025, focusing on industry distribution, OWASP API Top 10 alignment, and common abuse mechanics.

Patterns and Trends in API Breaches

INDUSTRY DISTRIBUTION

APIs get breached where the data and automation live

SECTORS/API BREACHES

Enterprise Software	9 (15%)
AI platforms and tooling	9 (15%)
Cybersecurity vendors	8 (13%)
SaaS providers	5 (8%)
Automotive	4 (7%)
Cloud services	4 (7%)
Other (media, telecom, retail, etc)	21 (35%)

The incident set clustered heavily in a few sectors.

Breaches concentrate in ecosystems with dense integrations, valuable data, and high levels of automation. If your business model depends on APIs, your risk model does too.

If your business model depends on APIs, your risk model does too.

OWASP API TOP 10 MAPPING

Identity and trust failures dominate

OWASP CATEGORY/API BREACHES	
API2 Broken Authentication	31 (52%)
API10 Unsafe Consumption of APIs	16 (27%)
API7 Security Misconfiguration	8 (13%)
API9 Improper Inventory Management	7 (12%)
API1 Broken Object Level Authorization	6 (10%)

Identity remains the fastest path from internet exposure to breach, which aligns closely with the OWASP vulnerability distribution earlier in this report. Categories like Broken Authentication, Unsafe Consumption of APIs, and Security Misconfiguration dominate both vulnerability disclosures and breach outcomes. Unsafe trust in third-party APIs and integrations consistently turns those design-time weaknesses into runtime incidents, expanding breach blast radius once attackers obtain even limited access.

API THREATSTATS TOP 10 MAPPING

Breach mechanics tell a clearer story

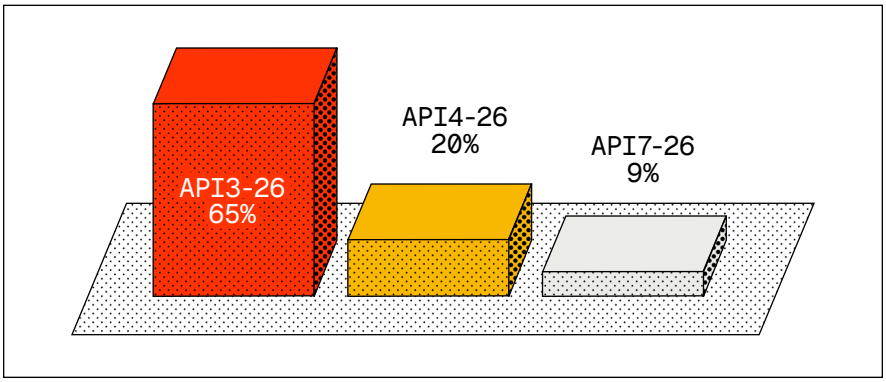

Breach analysis mapped to ThreatStats highlights identity and exposure as primary failure modes

API3-26: Authentication Flaws
API4-26: API Leaks
API7-26: Insecure Resource Consumption

When breaches are mapped directly to the dominant API abuse mechanism rather than disclosure categories, a sharper pattern emerges:

THREATSTATS CATEGORY/BREACHES	
API3-26 Authentication Flaws	30 (65%)
API4-26 API Leaks	9 (20%)
API7-26 Insecure Resource Consumption	4 (9%)
API1-26 Injections	2 (4%)
API10-26 SSRF	1 (2%)

This ThreatStats-based view shows that most API breaches are driven by failures in identity handling and exposure of sensitive interfaces, with automation and resource abuse acting as secondary accelerants rather than primary breach triggers.

API ThreatStats Top 10 Breaches for 2025

1 700CREDIT/FINTECH	2 QANTAS/AIRLINE	3 SALESLOFT DRIFT/SOFTWARE
700Credit	QANTAS	Salesloft.
DAMAGE: 5.6M VICTIMS[1] OWASP: API10	DAMAGE: 6M RECORDS STOLEN[2] OWASP: API2	DAMAGE: UNDISCLOSED[3] OWASP: API2, API10
WHAT: A third-party API integration exposed consumer credit data due to unsafe trust and insufficient controls on downstream API access. **LESSON:** Third-party API trust defines breach blast radius.	**WHAT:** Attackers abused airline APIs with weak authentication controls to access customer records at scale. **LESSON:** Consumer-scale APIs plus weak authentication produce predictable outcomes.	**WHAT:** Compromised OAuth tokens were used to access connected SaaS APIs and downstream customer data. **LESSON:** OAuth tokens function as portable admin credentials when scope and lifetime are loose.

4 SWISSBORG/ CRYPTOCURRENCY	5 HOSTING PROVIDERS MPC/AI	6 ANTHROPIC/CLAUDE APIS
SwissBorg		Claude
DAMAGE: $41M STOLEN4 OWASP: API2	DAMAGE: 3,000+ SERVERS5 OWASP: API2, API10	DAMAGE: PLATFORM EXPOSED6 OWASP: API2, API10
WHAT: Stolen credentials and API access were used to initiate unauthorized transactions and drain user accounts. **LESSON:** Fintech attackers optimize for speed, identity compromise first, loss second.	**WHAT:** An API-reachable path traversal flaw exposed Model Context Protocol servers used by agentic AI tooling. **LESSON:** Agentic AI infrastructure multiplies risk when exposed APIs lack runtime controls.	**WHAT:** API access control weaknesses enabled unintended access to AI-related data or functionality. **LESSON:** AI APIs frequently sit in the middle of sensitive enterprise workflows.

7 LANGSMITH/AI TOOLING

INTERNAL SERVICE EXPOSURE[7]
OWASP: API2, API10

WHAT: Internal APIs were exposed without sufficient authentication, revealing prompts, traces, or operational data.

LESSON: Dev and observability APIs contain attacker-grade data.

8 STEAM/GAMING PLATFORM

89 MILLION RECORDS[8]
OWASP: API2

WHAT: Automated abuse of platform APIs enabled large-scale scraping or account-level activity against users.

LESSON: Massive user bases attract relentless automated abuse.

9 ORACLE CLOUD/CLOUD SERVICES

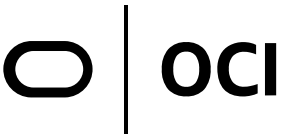

API-REACHABLE VULNERABILITY[9]
OWASP: API2, API10

WHAT: An internet-reachable cloud management API vulnerability allowed unauthorized access to control-plane functions.

LESSON: Control-plane APIs remain keys to the kingdom.

10 CONDÉ NAST (WIRED)

CONDÉ NAST

BOLA AND SCRAPING ABUSE[10]
OWASP: API1

WHAT: Broken object-level authorization allowed attackers to enumerate and scrape content via public APIs.

LESSON: Object-level authorization failures remain easy to ship and easy to exploit.

[1] Massive data breach sees credit card details of over 5.6 million victims leaked. [2] Australia's Qantas says 6 million customer accounts accessed in cyber hack. [3] Drift/Salesforce Security Update August 26, 2025 at 6:15 PM. [4] SwissBorg hacked for $41M SOL after third-party API compromise. [5] From Path Traversal to Supply Chain Compromise: Breaking MCP Server Hosting [6] Prompt Injection Variant Lets Hackers Exfiltrate Data from Claude APIs [7] LangSmith Bug Could Expose OpenAI Keys and User Data via Malicious Agents [8] Massive Alleged Steam Data Breach Results in Over 89 Million Records for Sale [9] Oracle Cloud Breach Exploiting CVE-2021-35587: How to Protect Your Organization [10] Santa's Naughty List? Wired Breach Gifts 2.3M Subscriber Records to the Dark Web

Breach Conclusions

The 2025 breach pattern was consistent and unsentimental:

1. IDENTITY FAILURES WERE THE PRIMARY BREACH ACCELERANT.
2. UNSAFE API CONSUMPTION EXPANDED BREACH IMPACT.
3. AUTOMATION WAS THE ATTACKER'S DEFAULT MODE OF OPERATION.

For defenders, the implications are clear: inventory APIs aggressively, treat tokens and secrets as production-critical assets, and design controls that assume abuse at scale. This matters even more as AI adoption accelerates. In 2025, AI platforms and tooling accounted for 15% of all API-related breaches, matching traditional software as the single largest industry category. These incidents followed the same mechanics seen elsewhere in the dataset, stolen tokens, exposed endpoints, and unsafe integration trust, but with higher potential impact because AI APIs often sit directly in the path of sensitive data, automation, and decision-making. APIs do not just enable the business, they define how it can be compromised.

Key Takeaways

For Security Practitioners

These findings translate into concrete changes practitioners can make in how APIs are built, tested, and defended.

DESIGN FOR ABUSE, NOT JUST CORRECTNESS. Across attacks, exploits, and breaches, the dominant failure modes are authentication gaps, exposed interfaces, and automation-friendly endpoints. Controls that only detect patterns over time will miss the single-request, low-friction abuse that dominates API risk.

TREAT IDENTITY AS THE PRIMARY ATTACK SURFACE. In breach analysis, 65% of incidents mapped to Authentication Flaws when classified by dominant mechanism. Short-lived tokens, strict scope enforcement, and continuous monitoring of token usage matter more than adding new perimeter controls.

ASSUME EVERY EXPOSED API WILL BE DISCOVERED AND SCRIPTED. Vulnerability data shows nearly all API flaws are remotely exploitable, and attack telemetry shows resource abuse and enumeration scale quickly once endpoints are found. Rate limiting, behavioral detection, and session-aware controls should be default, not optional.

SECURE AGENTIC APIS AS PRODUCTION INFRASTRUCTURE. MCP-related vulnerabilities and breaches show that APIs controlling tools and actions amplify impact. Over-permissioned agents and weak runtime enforcement turn ordinary API bugs into autonomous execution paths.

INVENTORY IS A SECURITY CONTROL. Exposed admin endpoints, shadow APIs, and forgotten integrations repeatedly appear in breach data. Continuous API discovery and classification are prerequisites for effective protection.

For CISOs and Security Leaders

At the leadership level, the data points to where investment and accountability have the highest return.

API RISK IS BUSINESS RISK. APIs dominate vulnerabilities, exploits, and breaches because they sit directly on revenue flows, customer data, and automation. Security programs that treat APIs as a niche concern will continue to absorb outsized incidents.

AI SECURITY FAILURES ARE OVERWHELMINGLY API-DRIVEN. 36% of AI vulnerabilities and 36% of AI-related exploited vulnerabilities also involve APIs. Investments in AI security that do not include strong API controls will leave material gaps.

IDENTITY FAILURES DRIVE BREACH OUTCOMES. When breaches are classified by how they actually occurred, authentication flaws account for nearly two-thirds of incidents. This argues for executive focus on identity governance, token lifecycle management, and third-party access, not just vulnerability counts.

AUTONOMY INCREASES IMPACT. MCP illustrates where risk is heading: APIs that act on behalf of systems rather than users. As automation increases, control failures delegate power to attackers. This demands governance over what APIs and agents are allowed to do, not just who can call them.

MEASURE WHAT ATTACKERS EXPLOIT, NOT JUST WHAT IS DISCLOSED. The ThreatStats Top 10, KEV analysis, and breach mapping all converge on the same message: runtime behavior matters more than theoretical risk. Align metrics, tooling, and reporting to observed abuse, not checklists.

Taken together, the 2025 data shows that improving API security is not about chasing new attack classes. It is about fixing repeatable failures, identity, exposure, and abuse, before scale and automation turn them into incidents.

About Wallarm

Wallarm is the leading API security company, purpose-built to protect modern cloud-native and AI-driven architectures from today's most advanced threats. Our platform delivers complete visibility, intelligent threat detection, and real-time protection for all types of APIs like REST, GraphQL, gRPC, and increasingly, AI and Agent-based APIs.

As organizations adopt LLMs and autonomous agents, Wallarm helps secure the unique risks these interfaces introduce, including prompt injection, token abuse, and logic manipulation. By combining continuous API discovery, behavior-based analysis, and runtime policy enforcement, Wallarm enables security teams to protect complex API ecosystems including those powering AI without slowing innovation.

LEARN MORE

Website	wallarm.com
Blog	lab.wallarm.com
X (Twitter)	x.com/wallarm
LinkedIn	linkedin.com/company/wallarm
YouTube	youtube.com/@wallarmchannel
Explore Product	tour.playground.wallarm.com