wallarm

Q2’24

API ThreatStats™ Report

M&A Mayhem

API breaches in TestRail, HelloSign, Duo, and Authy expose vulnerabilities introduced during mergers and acquisitions.

AI API Exploit Explosion

Tripled API vulnerabilities in Anything-llm and ZenML signal urgent security needs.

JWT Jumbles

Persistent API issues with JWT in Veeam, Lua-Resty-JWT, and Python-jose highlight authentication challenges.

www.wallarm.com

(415) 940-7077
188 King St. Unit 508, San Francisco, CA 94107

Introduction

Welcome to the API ThreatStats™ Q2’24 report. This comprehensive analysis of the evolving API security landscape is presented by Wallarm, with a focus on providing empirical data and actionable insights. Our team of researchers and security experts has meticulously gathered and analyzed data to ensure that our findings are grounded in evidence.

In Q2 2024, we observed a significant increase in AI API exploits, with the number of identified vulnerabilities tripling compared to the previous quarter. This trend, predicted in the Q1 report, underscores the growing importance of securing AI systems like Anything-llm and ZenML as they become more integrated into the digital ecosystem. The volume and complexity of these vulnerabilities highlights the need for robust security measures to protect these critical infrastructures.

The impact of mergers and acquisitions (M&A) on API data breaches is another key finding in this quarter’s report. Notable incidents involving TestRail (Atlassian), HelloSign (Dropbox), Duo (Cisco), and Authy (Twilio) emphasize that vulnerabilities introduced during M&A activities introduce impactful risk to the acquiring organizations. These cases illustrate the critical need for thorough security assessments, due diligence, and the integration of stringent security protocols during such transitions.

The continued misuse of JSON Web Tokens (JWT) across various applications remains a significant concern. Despite JWT's widespread adoption as a standard for securing API communications, proper implementation remains a challenge. This quarter's data includes significant vulnerabilities in enterprise solutions like Veeam Recovery Orchestrator and open-source frameworks such as Lua-Resty-JWT and Python-jose, highlighting JWT as a major authentication issue that spans both proprietary and community-driven software.

Vulnerabilities discovered in popular platforms like Grafana, despite its strong security focus and active community, were particularly surprising. These incidents underscore the necessity for continuous monitoring and proactive security practices.

This report is uniquely structured to provide tailored insights for different audiences. Based on feedback from previous reports, the conclusions are divided into two parts: one for CISOs and another for security practitioners. This approach addresses both high-level strategic concerns and detailed technical issues, ensuring that the findings are relevant and useful to all stakeholders.

Wallarm remains committed to providing precise data points to score and measure API threats, avoiding reliance on assumptions. The goal is to empower organizations with the knowledge and tools necessary to safeguard against the ever-evolving landscape of cybersecurity threats.

Readers are invited to delve into the detailed findings of this report to gain a comprehensive understanding of the current threat landscape. Whether responsible for strategic decisions or hands-on technical implementations, this report is designed to equip professionals with the insights needed to enhance their organization’s security posture.

Thank you for your continued dedication to security.

Sincerely,

Ivan Novikov

Ivan Novikov
CEO, Wallarm

API Related Data Breaches in Q2 2024

Date:	Company:	Users Affected:	Data Exposed:	Industry:	
April 16, 2024 Discovered: April 1, 2024	DUO	170 000	SMS logs, phone numbers, metadata	Enterprise	API Leak
April 22, 2024 Discovered: March 27, 2024	digiyatra	1 740 000	Aadhaar details, flight history, personal preferences	Aviation	
April 29, 2024 Discovered: April 24, 2024	HELLOSIGN	100 000	Emails, usernames, phone numbers, hashed passwords, API keys, OAuth tokens	SaaS	
June 25, 2024 Discovered: May 16, 2024	rabbit™	100 000	All responses given by R1 devices,	Consumer Electronics	
June 2024 Discovered: June 2024	TestRail	11 000	Potential compromise of user API tokens	SaaS	
June 4, 2024 Discovered: June 2024	TECHINASIA	230 000	Email addresses, names, user IDs, user roles	Media	Broken Access Control
June 21, 2024 Discovered: September 2022	OPTUS	9 500 000	Customer information	Telco	
April 28, 2024 Discovered: March 2024	DELL™	49 000 000	Warranty info, service tags, customer data	Enterprise	API Abuse
July 3, 2024 Discovered: July 1, 2024	twilio	33 000 000	Phone numbers, carriers, metadata	SaaS	Broken Authentication

The second quarter of 2024 saw a significant surge in data breaches across various industries, revealing alarming weaknesses in API security. These incidents highlight the urgent need for improved cybersecurity measures and offer valuable insights into the vulnerabilities that different sectors face.

The Most Affected Sectors

The enterprise sector was hit hardest by API vulnerabilities, affecting 49 million users. Dell's partner portal API was exploited, allowing data scraping from 49 million users. Cisco Duo Security also reported a breach through its third-party telephony provider, exposing SMS logs and metadata for 170,000 Duo users. These incidents highlight the critical need for stringent security protocols and regular reviews of API security measures in large corporations. Many breaches stem from M&A activities where acquired companies impact the acquirer's API security.

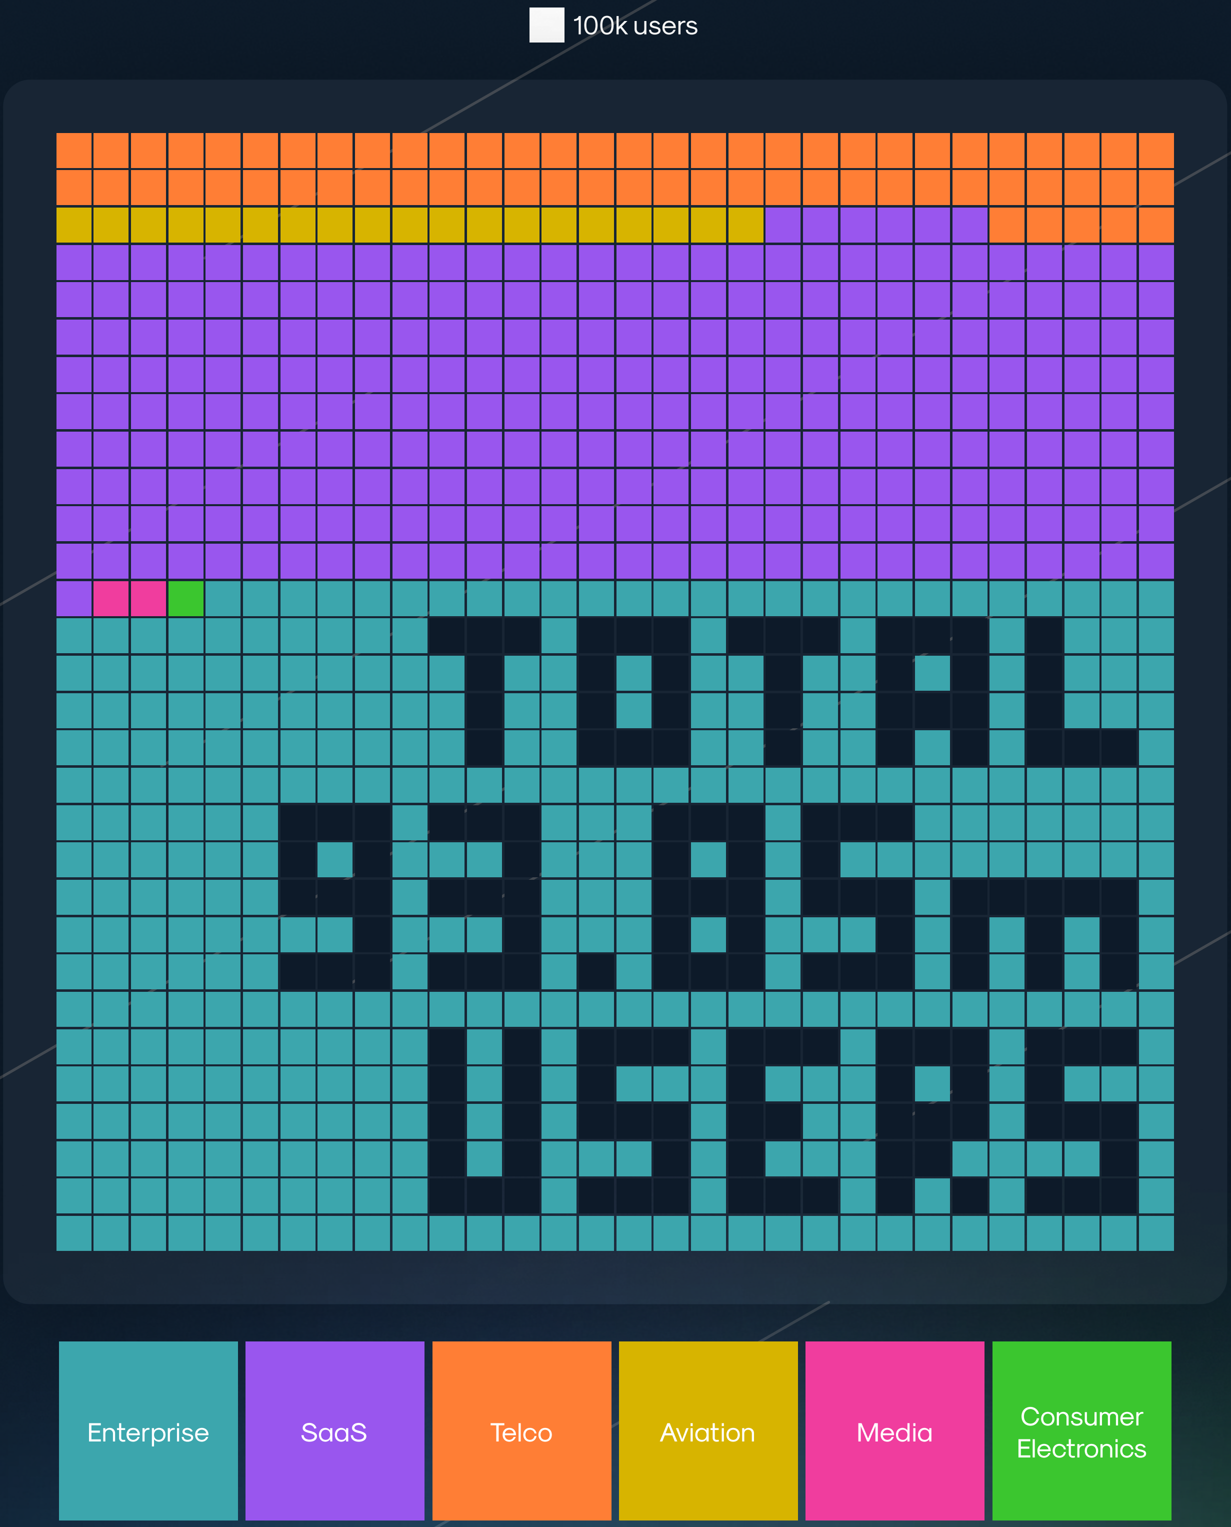

SaaS providers followed, with three major breaches impacting 33 million users. Dropbox HelloSign, Twilio, and TestRail/Atlassian suffered from API vulnerabilities, underscoring the sector's reliance on secure API integrations. Dropbox HelloSign exposed emails, usernames, phone numbers, hashed passwords, API keys, and OAuth tokens. Twilio's breach due to broken authentication compromised phone numbers and metadata of 33 million users. TestRail/Atlassian faced an API leak, potentially compromising user API tokens.

Telecommunications, represented by Optus, experienced a breach due to broken access control, affecting 9.5 million users. This incident highlights the necessity for continuous monitoring and maintenance of security controls to prevent unauthorized access.

The aviation industry, through Digi Yatra, faced an API leak that exposed the Aadhaar details, flight history, and personal preferences of thousands of Indian frequent flyers. This breach illustrates the risks associated with integrating modern digital solutions in traditional sectors and the need for stringent security checks.

Media and consumer electronics sectors were not spared either. Tech in Asia, a media company, suffered a breach due to broken access control, compromising the personal information of 230,000 users. Rabbit Inc., representing consumer electronics, faced an API leak that put hundreds of R1 device users at risk by exposing all responses given by the devices along with personal information. These breaches reveal that vulnerabilities can exist across various platforms, from content management to IoT devices.

Major API Threats Types

API leaks were the most common type of breach, accounting for over half of the incidents. This indicates a widespread issue with how APIs are secured and managed. The frequency of API leaks, particularly among SaaS providers and enterprise sectors, highlights the urgent need for comprehensive security measures, including regular audits, stringent access controls, and real-time monitoring. It's worth noting that we highlighted API Leaks as a growing threat in our annual API ThreatStats™ report covering all of 2023.

Another significant type of breach was broken access control, affecting both the telecommunications and media sectors. This type of vulnerability underscores the importance of continuous monitoring and maintenance of security controls to prevent unauthorized access and data breaches.

The analysis also revealed that broken authentication, although less frequent, had a substantial impact, as seen in Twilio's breach. And the API Abuse incident at Dell highlights the need for mandatory granular endpoints controls such as key usage, rate limiting, and data excursion.

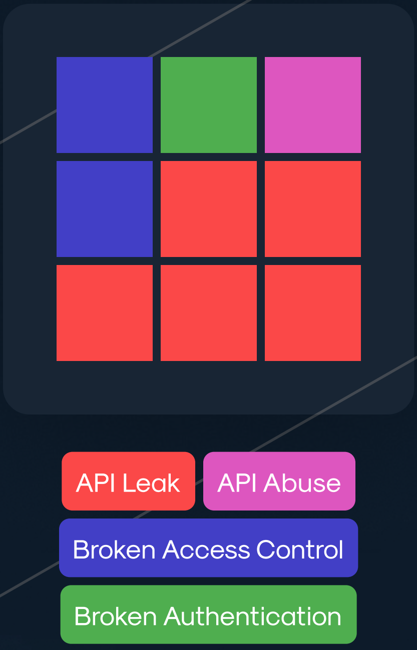

OPTUS

Date of Breach:
Discovered: September 2022

#1 Optus Data Breach

API Vulnerability Details:

- A coding error in 2018 broke API access controls on both main (www.optus.com.au) and target (api.optus.com.au) domains.
- Error fixed on the main domain in 2021, but left unchecked on the target domain.
- Target domain remained online and vulnerable.

Incident Details:

- Simple attack using trial and error.
- Accessed customer information via Target APIs.
- Exploited due to poor access control and maintenance.

Regulatory Actions:

- Australia's Communications and Media Authority (ACMA) pursuing civil penalties.
- Optus and parent company Singtel involved in legal defense.
- Court filing happened in Q2 2024, detailing the incident and legal proceedings.

Court Orders (June 2024):

- Optus to produce a final report by Deloitte by June 21, 2024.
- Confidentiality agreements apply to the report.
- Compliance date extended to September 13, 2024.
- Case management hearing scheduled for September 13, 2024.
- Costs are reserved, with liberty to apply on two days' notice.

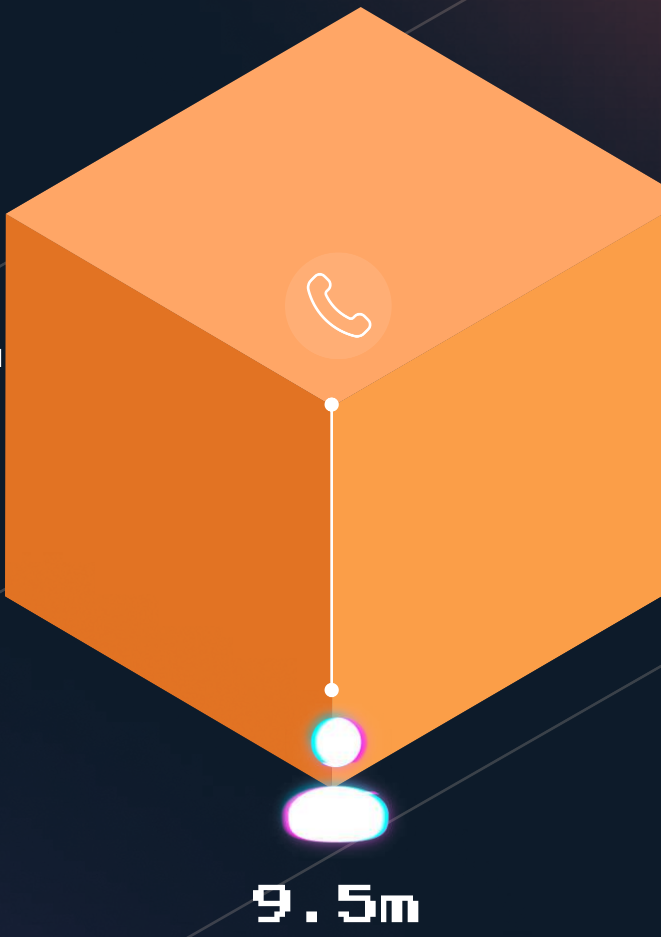

Key Points:

- Breach blamed on overlooked API error.
- Poor access control and maintenance led to prolonged vulnerability.
- Attack did not require sophisticated skills.

Date of Breach:
Discovered: March 2024

#2 Dell Data Breach

API Vulnerability Details:

- Partner portal API abused by a threat actor posing as fake companies.
- No verification for partner registration.
- Lack of rate limiting allowed massive data scraping.

Incident Details:

- Threat actor Menelik created multiple fake accounts.
- Used a program to generate 7-digit service tags and scraped data.
- Generated 5,000 requests per minute over three weeks.
- Stolen data includes warranty information, service tags, customer names, installed locations, customer numbers, and order numbers.

Stolen Customer Records by Hardware:

Monitors:	**22,406,133**
Alienware Notebooks:	**447,315**
Chromebooks:	**198,713**
Inspiron Notebooks:	**11,257,567**
Inspiron Desktops:	**1,731,767**
Latitude Laptops:	**4,130,510**
Optiplex:	**5,177,626**
Poweredge:	**783,575**
Precision Desktops:	**798,018**
Precision Notebooks:	**486,244**
Vostro Notebooks:	**148,087**
Vostro Desktops:	**37,427**
XPS Notebooks:	**1,045,302**
XPS/Alienware Desktops:	**399,695**

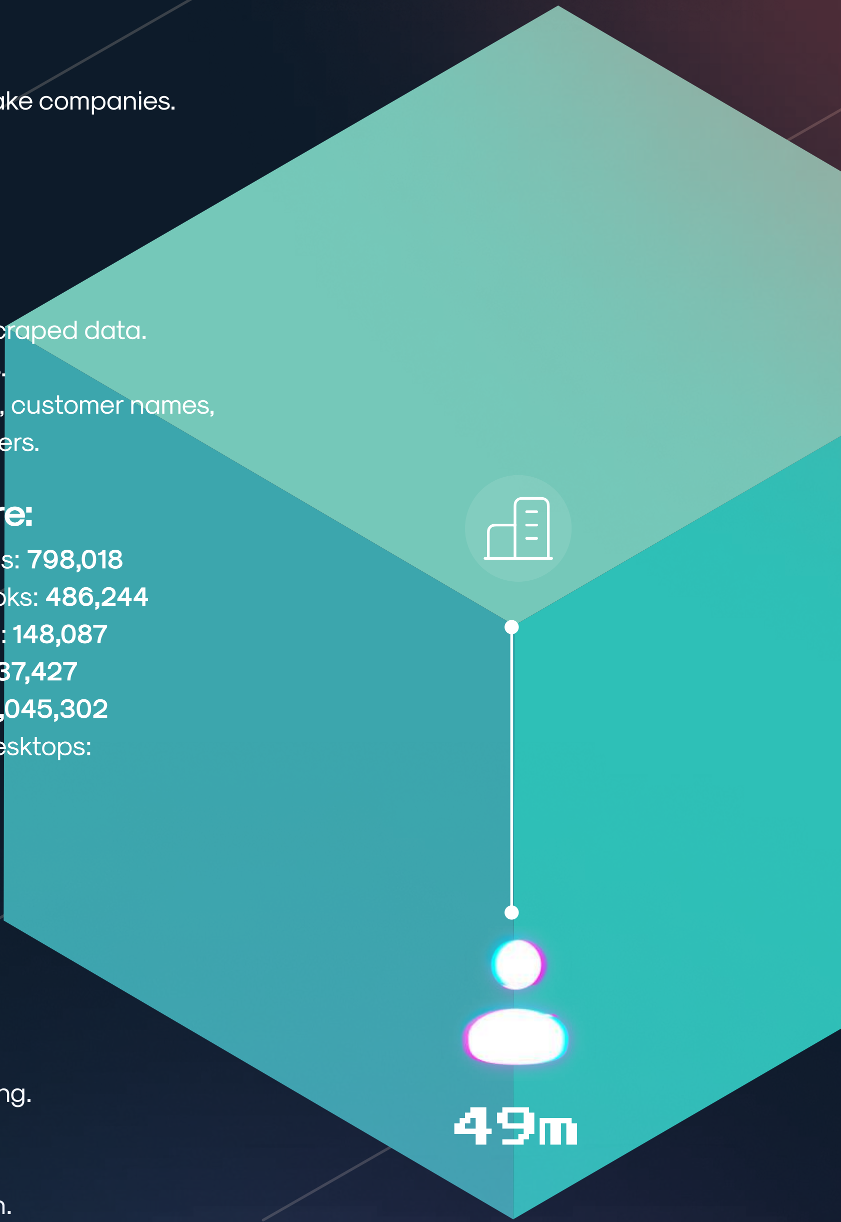

Regulatory Actions:

- Law enforcement investigation initiated.
- Dell engaged a third-party forensics firm.

Key Points:

- Breach due to lack of proper API security and rate limiting.
- Exploited partner portal's easy registration process.
- Massive data scraping over several weeks.
- Dell's delayed response to the threat actor's notification.

Timeline:

- Breach activity started: March 2024.
- Threat actor reported bug to Dell: April 12th and 14th, 2024.
- Data listed for sale: April 28, 2024.
- Dell notified customers: May 2024.
- Court filing occurred in Q2 2024.

HELLOSIGN

Date of Breach:
Discovered: April 24, 2024

#3 Dropbox Data Breach Summary

API Vulnerability Details:

- Unauthorized access to Dropbox Sign production environment.
- Accessed data: emails, usernames, phone numbers, hashed passwords, API keys, OAuth tokens, and multi-factor authentication information.
- No access to user account contents, agreements, templates, or payment information.

Incident Details:

- Threat actor accessed data via compromised API keys and authentication tokens.
- Incident isolated to Dropbox Sign infrastructure.
- Immediate response with cybersecurity measures and forensic investigation.

Regulatory Actions:

- Notified and cooperating with law enforcement.
- Notifying regulatory authorities and affected users.

Timeline:

- Breach detected: April 24, 2024.
- Public disclosure: May 1, 2024.
- Incident report filed: April 29, 2024.

Key Points:

- Breach limited to Dropbox Sign; no impact on other Dropbox products.
- Ongoing investigation and mitigation efforts.
- Potential risks: litigation, customer behavior changes, regulatory scrutiny.
- Current understanding suggests no material impact on business operations or financial condition.

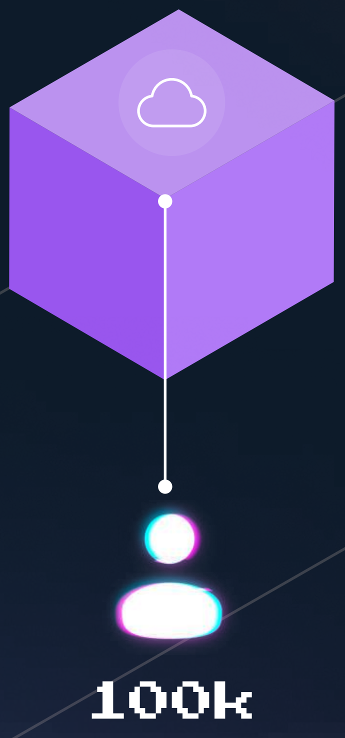

Links:

https://www.sec.gov/Archives/edgar/data/1467623/000146762324000024/dbx-20240429.htm

Date of Breach:
Discovered: July 1, 2024

#4 Twilio Data Breach Summary

API Vulnerability Details:

- Unauthenticated endpoint in Authy's API allowed threat actors to identify phone numbers associated with Authy accounts.
- Endpoint now secured; unauthenticated requests no longer allowed.

Incident Details:

- Threat actor ShinyHunters accessed phone numbers by inputting a massive list into Authy's unsecured API.
- No evidence of access to Twilio's main systems or other sensitive data.
- Possible risk of phishing and smishing attacks using the obtained phone numbers.

Regulatory Actions:

- Twilio notified law enforcement and relevant regulatory authorities.
- Advised users to update Authy apps for latest security updates and remain vigilant against phishing attacks.

Timeline:

- Unauthorized access: Discovered in July 2024.
- Threat actor posted stolen phone numbers: Published last week of June 2024.
- Public disclosure: July 3, 2024.

Key Points:

- Breach involved enumeration of phone numbers via unsecured API.
- Immediate response to secure endpoint and mitigate risk.
- Previous 2022 breach related to phishing campaign targeting employees and customers.
- Ongoing investigation and user notification to prevent further risks.

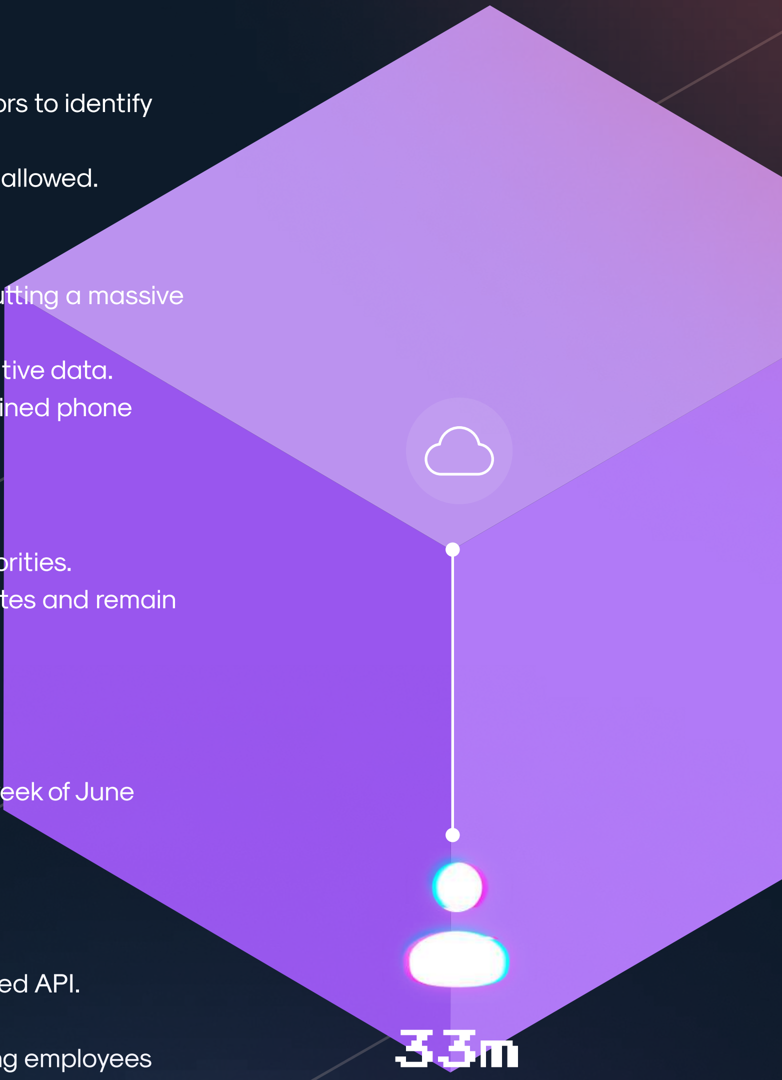

Date of Breach:
Discovered: May 16, 2024

#5 Rabbit Inc. Data Breach Summary

API Vulnerability Details:

- Hardcoded API keys discovered in Rabbit's codebase.
- Services affected:
  - ElevenLabs (Text-to-Speech)
  - Azure (Old Speech-to-Text)
  - Yelp (Review Lookups)
  - Google Maps (Location Lookups)

Incident Details:

- Threat actor "evil-rabbitude" accessed Rabbit codebase.
- API keys allowed:
  - Reading all R1 responses, including personal information.
  - Bricking all R1 devices.
  - Altering R1 responses.
  - Replacing R1 voices.
- ElevenLabs API key provided full privileges for text-to-speech manipulation.
- Rabbit Inc. aware of the issue for a month without rotating the API keys.

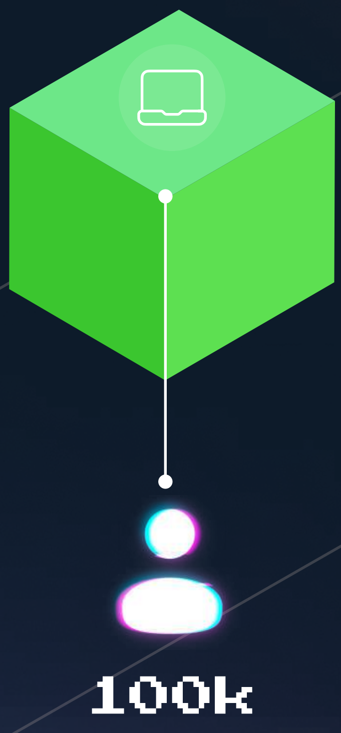

Key Exploits:

- Access to history of all past text-to-speech messages.
- Ability to change voices and add custom text replacements.
- Capability to delete voices and crash RabbitOS backend, rendering R1 devices useless.

Key Points:

- Breach due to hardcoded API keys in codebase.
- Significant security risk for R1 users due to potential exposure of personal information.
- Rabbit Inc.'s inaction highlights poor security practices.
- Ongoing concerns for consumer data safety and device integrity.

Rabbit Inc. Response:

- Internal confirmation of awareness of the API key leak.
- No action taken to rotate or invalidate the API keys.
- Consumers advised to unlink RabbitHole connections for security.

Timeline:

- API keys discovered: May 16, 2024
- Public disclosure by xyzeva: June 25, 2024

Date of Breach:
Discovered: March 27, 2024

#6 Digi Yatra Data Breach Summary

API Vulnerability Details:

- API endpoint of old Digi Yatra app communicated with DataEvolve's AWS servers.
- New API endpoint established to secure data.

Incident Details:

- DigiEvolve, the app maker, was dropped by Digi Yatra Foundation due to compromised data.
- Previous app versions sent passenger data to DataEvolve's servers.
- Users instructed to uninstall the old app and install a new version with a new API endpoint.
- Privacy concerns raised about data storage policies and security measures.

Regulatory Actions:

- Civil Aviation Ministry's response on January 24, 2024, claimed data was stored on passengers' mobiles.
- Security audits included penetration testing, code reviews, and network security assessments but lacked background checks on the app maker.

Key Points:

- Hardcoded API keys and poor data security practices by DataEvolve.
- Significant privacy risks for users due to potential misuse of personal data.
- Transition to a new app and API endpoint to secure user data.
- Previous legal issues with DataEvolve raised concerns about their reliability.
- Emphasis on improved security practices and user vigilance against potential data misuse.

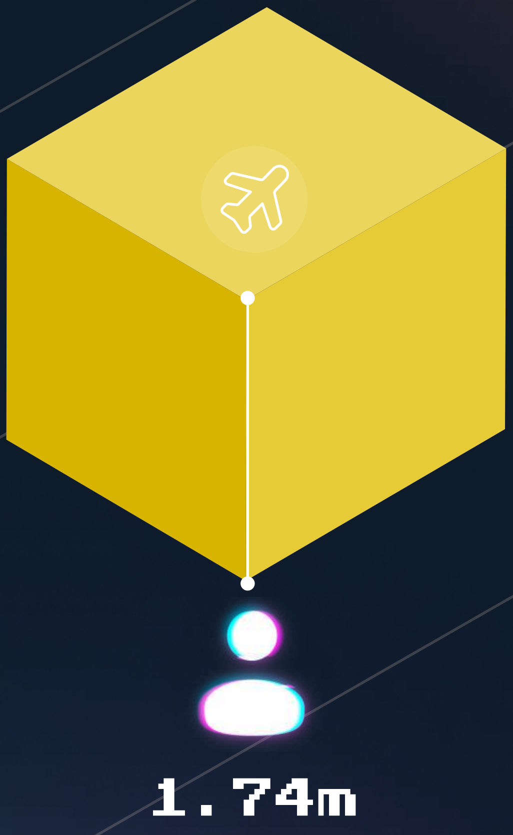

Timeline:

- App compromise discovered: March 27, 2024.
- Public disclosure and user instructions: April 22, 2024.

Date of Breach:
Discovered: April 1, 2024

#7 Cisco Duo Data Breach Summary

API Vulnerability Details:

- Breach occurred at a third-party telephony provider.
- Attackers accessed systems via stolen credentials obtained through phishing.
- Stolen data included logs of MFA SMS messages.

Incident Details:

- Attackers downloaded SMS logs for authentication messages sent between March 1 and March 31, 2024.
- Data breach puts customers at risk of social engineering attacks, credential theft, and financial fraud.

Regulatory Actions:

- Cisco notified affected customers on April 16, 2024.
- Third-party provider invalidated the phished employee's credentials.
- Cisco offered to provide copies of the compromised message logs to affected customers upon request.

Key Points:

- Attack targeted a third-party service provider critical to Cisco Duo's MFA services.
- Stolen SMS logs can be used for phishing attacks and other fraudulent activities.
- Highlights the importance of third-party security and the risks associated with API suppliers.
- Cisco Duo has over 100,000 users across 98 countries, amplifying the breach's potential impact.

Expert Opinions:

- Jim Routh, Saviynt: Breach makes it easier for attackers to target users with phishing lures.
- Jeff Margolies, Saviynt: Emphasizes trend of attacks on identity security providers and third-party vulnerabilities.
- Jamie Beckland, API Context: Stresses the need for real-time tracking of API suppliers for rapid security responses.

Timeline:

- Breach occurred: April 1, 2024.
- Notification to customers: April 16, 2024

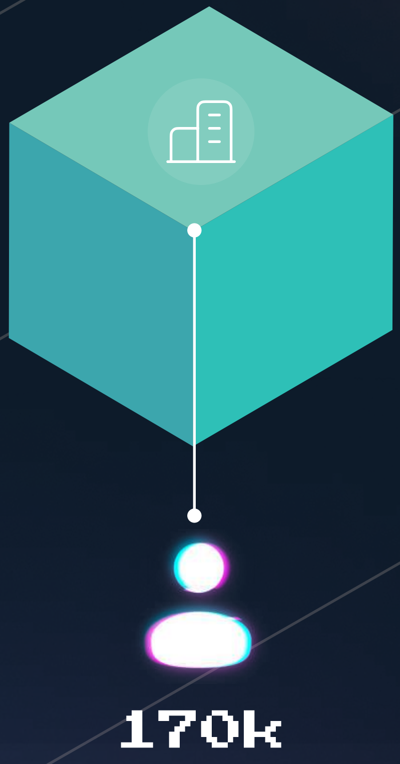

TECHINASIA

#8 Tech in Asia Data Breach Summary

Date of Breach:
Discovered: June 2024

API Vulnerability Details:

- Exploited vulnerabilities within Tech in Asia's API.
- Unauthorized access to user data and internal services.

Incident Details:

- Threat actor Sanggiero published the leaked data on Breach Forums.
- Data includes sensitive information like email addresses and full names, potentially leading to identity theft and targeted attacks.
- User data: User ID, Tech in Asia ID, email address, user roles, full name, display name, registration date, avatar URL, author URL.

Regulatory Actions:

- Awaiting official statement and response from Tech in Asia.

Key Points:

- Breach exploited API vulnerabilities.
- Significant privacy risks for users, including potential identity theft and phishing attacks.
- Users advised to change passwords, beware of phishing attempts, and monitor accounts for unusual activity.

User Precautions:

- **Change Password:** Update passwords on Tech in Asia and any other accounts using the same credentials.
- **Beware of Phishing:** Be cautious of emails requesting personal information or containing suspicious links.
- **Monitor Accounts:** Stay alert for unusual activity on Tech in Asia and linked accounts.

Response and Mitigation:

- Tech in Asia is expected to release a statement outlining steps to safeguard user data and prevent future breaches.
- Users should stay informed about updates from Tech in Asia regarding the breach.

Timeline:

- Data breached: June 2024.
- Public disclosure: June 4, 2024.

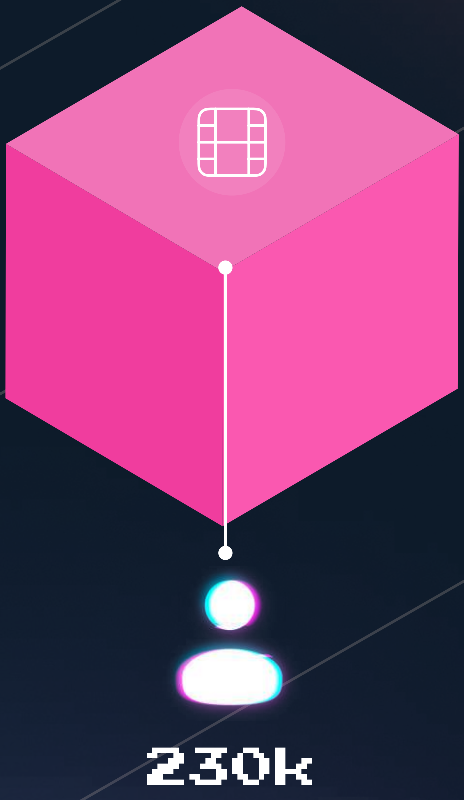

Date of Breach:
Discovered: June 2024

#9 Atlassian/TestRail Data Breach Summary

API Vulnerability Details:

- Unusual activity detected on other customers' instances.
- Suspected compromise of user API tokens associated with TestRail.

Incident Details:

- Atlassian notified users about the potential compromise and revoked API tokens as a precaution.
- Compromised tokens could allow access to Jira/Confluence integrations, exposing sensitive data.
- Users expressed concerns about phishing attacks using the compromised data.

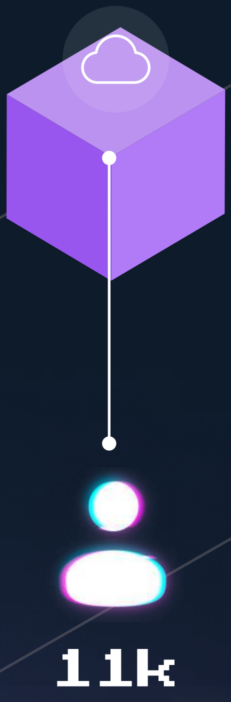

Regulatory Actions:

- TestRail and Atlassian working together to address the issue.
- TestRail prompted users to reset Jira integration settings.

Key Points:

- Breach potentially involved compromised API tokens for TestRail.
- Proactive revocation of tokens by Atlassian to mitigate risks.
- Lack of clear communication from TestRail causing user frustration.
- Users advised to reset integrations and be vigilant against phishing attacks.

User Precautions:

- **Change Integration Settings:** Reset Jira and other integrations associated with TestRail.
- **Beware of Phishing:** Be cautious of unexpected emails or messages requesting credentials or containing suspicious links.
- **Monitor Accounts:** Stay alert for unusual activity in connected services like Jira and Confluence.

Response and Mitigation:

- Users awaiting further updates and clear communication from TestRail.
- Emphasis on proactive security measures and user vigilance to prevent further exploitation.

Timeline:

- Unusual activity noticed: June 2024.
- Notification to users: June 2024.
- Ongoing communication and investigation by TestRail and Atlassian.

Lessons Learned: MnA and 3rd party APIs are on fire

Q2 2024 breaches highlight the need for robust API security, especially during mergers and acquisitions. Notable breaches include TestRail (Atlassian), HelloSign (Dropbox), Duo (Cisco), and Authy (Twilio), emphasizing vulnerabilities introduced through acquisitions. Here area few action items we suggest:

Security Audits and Due Diligence

Security Audits and Due Diligence: Essential during M&A to assess third-party APIs. Integrations must meet stringent security standards.

Comprehensive API Security

Regular audits, stringent access controls, and real-time monitoring are crucial. Adopt best practices to safeguard user data as part of the acquisition process.

Proactive Incident Response

Develop and test robust incident response plans. Quick mitigation of breaches is vital.

Cross-Industry Vigilance

API security is a universal concern. Continuous improvement in security practices is necessary to protect against evolving threats.

Emphasize thorough security assessments during M&A, prioritize API security measures, and maintain proactive incident response strategies. Protect sensitive data and prevent future breaches.

API CVE Analysis

We began our API vulnerability analysis by examining the Top-5 vendors and uncovered a widespread issue: API security vulnerabilities are present across a broad spectrum of domains. Last quarter in Q1, we highlighted AI APIs as a growing trend, and this quarter, the trend continues with two AI/LLM tools, Mintplex and Lunary, making it into the Top-5 vulnerable vendors list for all API-related CVEs.

Our findings include large enterprise vendors like Dell, which indicates significant risks within well-established corporations. Our analysis also extended to major open-source projects from the Apache Software Foundation, revealing that even widely-used community-driven software is not immune to security flaws. Additionally, popular DevOps tools such as Grafana and GitLab were found to have notable vulnerabilities, highlighting risks in tools that many development teams rely on daily.

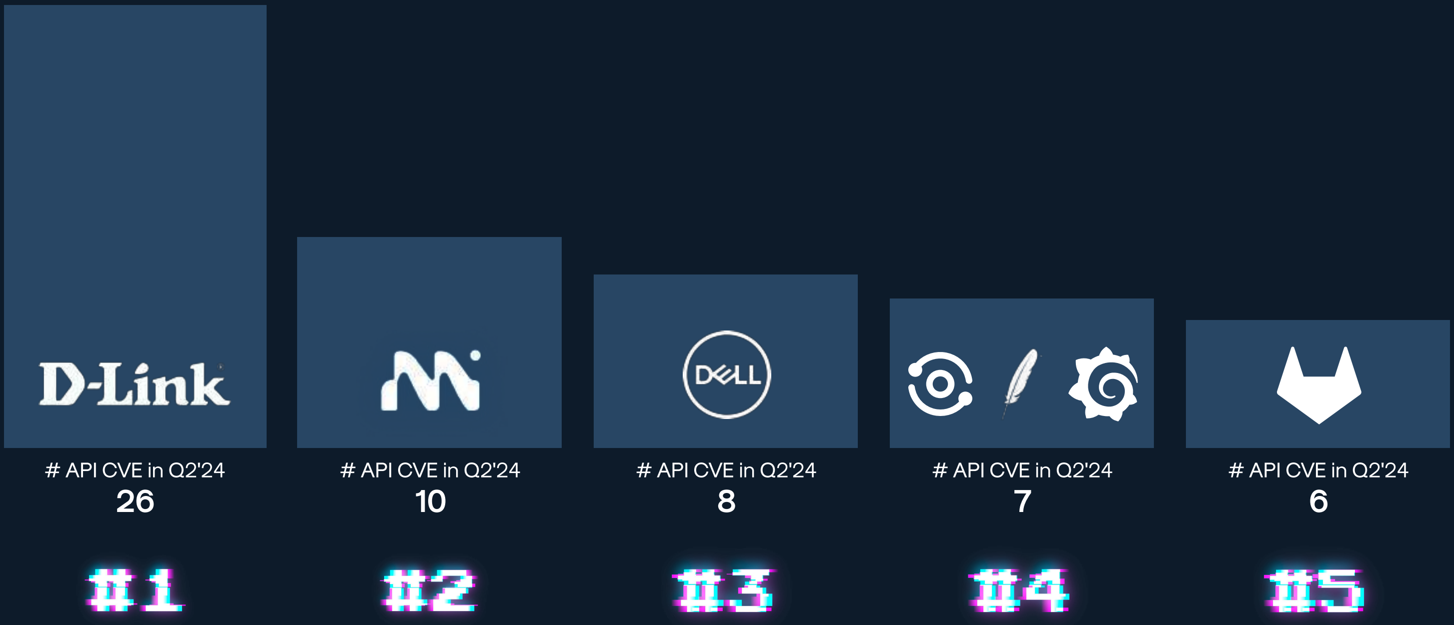

The conclusion from our comprehensive analysis is unmistakable: API security issues are omnipresent, striking hard across all sectors. This quarter's inclusion of AI/LLM tools among the most vulnerable vendors underscores the urgent need for enhanced security measures across both traditional and cutting-edge technologies.

Similar to this, the distribution of vulnerable product industries reveals a significant presence of AI/LLM tools across typical categories such as enterprise software, hardware, and DevOps. This further emphasizes the pervasive nature of API vulnerabilities and the necessity for comprehensive security measures across diverse sectors.

Our analysis showed that enterprise software products topped the list with 123 vulnerabilities, reflecting the extensive use and critical nature of these systems in business operations. DevOps tools followed with 69 vulnerabilities, underscoring the importance of securing tools that are integral to the development and deployment processes.

Enterprise hardware was not far behind, with 62 vulnerabilities detected, highlighting the risks within essential physical infrastructure. AI/LLM tools and frameworks accounted for 51 vulnerabilities, marking their growing significance and the emerging security challenges in this innovative field.

Development frameworks also showed substantial vulnerabilities, with 50 cases, pointing to the need for secure coding practices and robust framework security. Cloud services, despite their widespread adoption, reported 29 vulnerabilities, indicating that even well-established cloud environments require vigilant security practices.

This distribution clearly shows that API security is a critical concern across all major technology sectors, from traditional enterprise systems to cutting-edge AI and cloud services. It underscores the urgent need for a unified approach to API security that addresses the unique challenges within each category while maintaining a strong overall defense strategy.

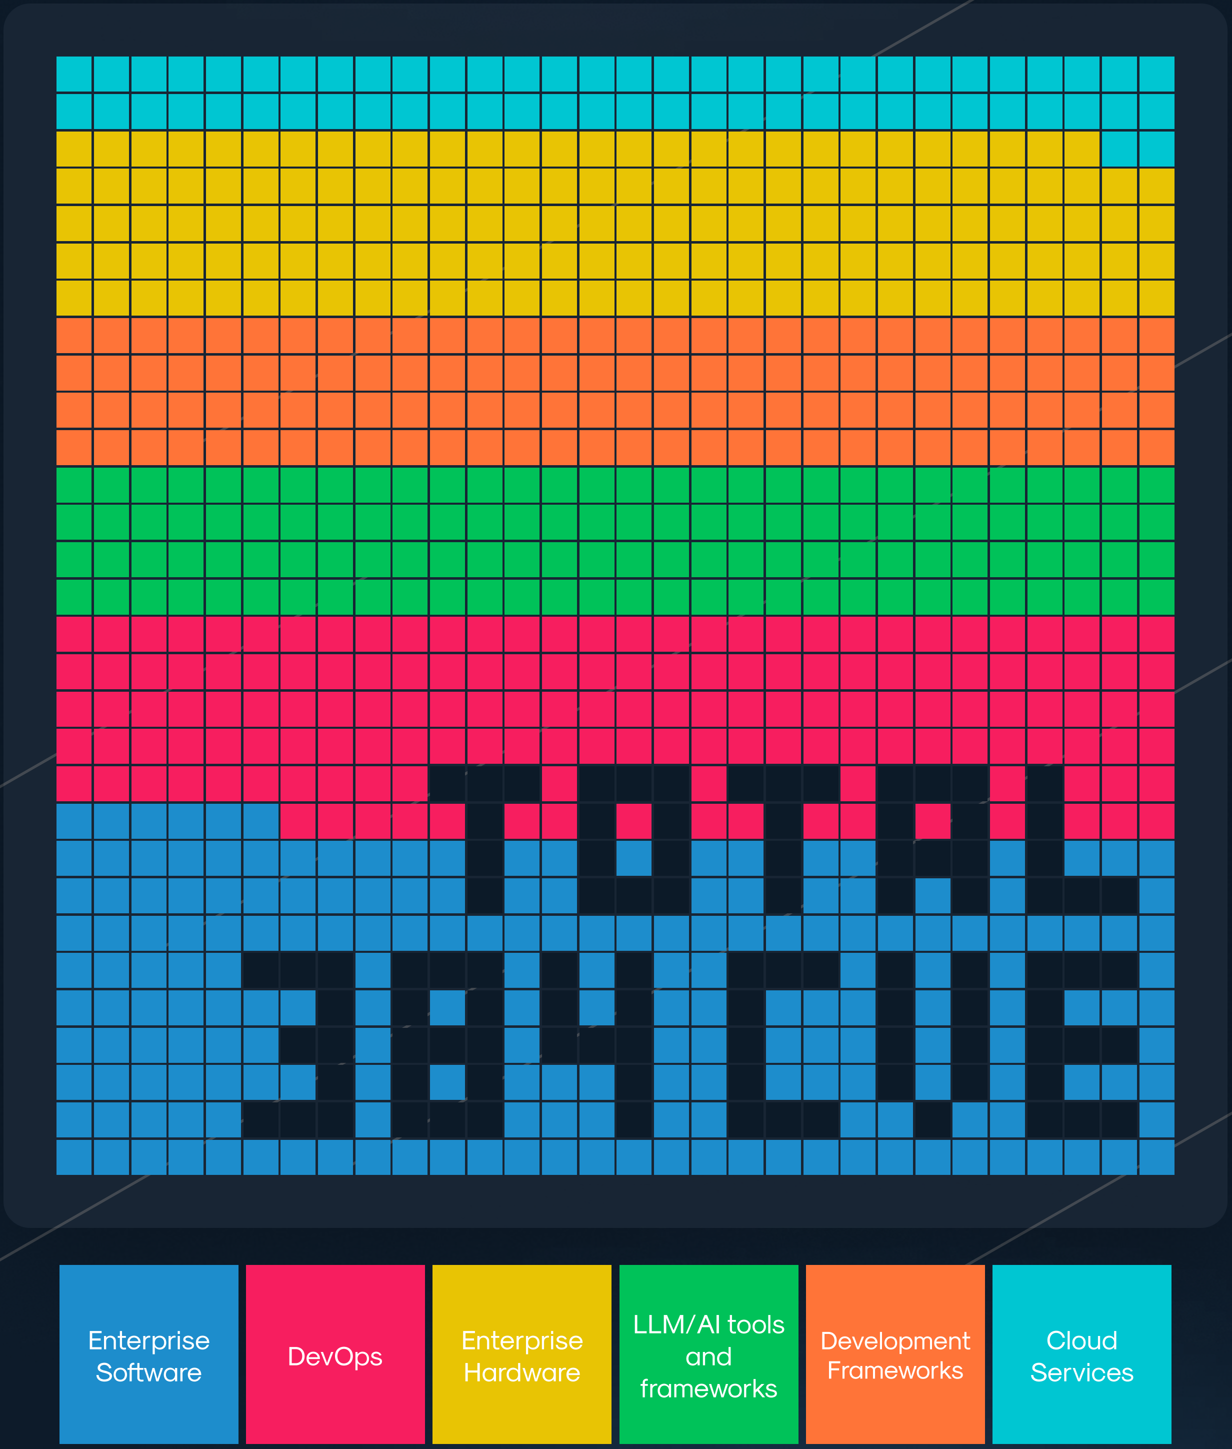

AI API Exploits Review

In our previous Q1 24 ThreatStats™ report, we highlighted AI API vulnerabilities as a rising star in the realm of cybersecurity. This trend has not only continued but has also intensified in Q2 2024. AI and machine learning (ML) systems are fundamentally driven by APIs, making these vulnerabilities critical to the integrity of AI-driven solutions. The interaction between AI models and the data they process is mediated by APIs, rendering them a critical attack vector. This chapter explores the various types of vulnerabilities discovered, their implications, and how they impact AI/ML/LLM (Large Language Models) APIs.

Remote Code Execution: The Apex Threat

Remote Code Execution (RCE) remains one of the most dangerous threats to AI APIs. In Q2 2024, several high-severity RCE vulnerabilities were identified. The Anything-llm API demonstrated critical weaknesses with an arbitrary file deletion vulnerability, scoring a 9.9, which allowed attackers to exploit path traversal flaws and execute commands remotely. Similarly, the Aim Web API's RCE flaw, scoring a severe 9.8, highlighted the potential for attackers to inject and execute arbitrary code through improperly sanitized input parameters. These vulnerabilities can allow adversaries to control the underlying infrastructure, manipulate data, disrupt AI operations, and even create backdoors for future attacks.

Path Traversal and Unauthorized Access

Path traversal vulnerabilities pose a significant threat to data integrity and confidentiality. The Anything-llm API's arbitrary file deletion issue, which exploited directory traversal flaws, enabled attackers to access and manipulate sensitive files outside the intended directories. This vulnerability scored a 9.9, illustrating the high risk associated with inadequate input validation. Similarly, ZenML's directory traversal vulnerability allowed unauthorized access to critical configuration files, scoring another 9.9. The improper access control in the allData API further underscores the risk of inadequate permission checks, where attackers can bypass authentication mechanisms to gain unauthorized access to sensitive data, scoring a 7.5.

Mass Assignment and Privilege Escalation

Mass assignment vulnerabilities, particularly prevalent in AnythingLLm, illustrate the dangers of improper user input handling. Attackers can manipulate JSON payloads to assign unauthorized attributes, leading to privilege escalation. The vulnerability in account creation from invitations, scoring a high 9.1, demonstrated how attackers could elevate their privileges to administrator level by injecting malicious parameters. This type of vulnerability highlights the importance of stringent input validation and access control mechanisms to prevent unauthorized privilege escalation. Additionally, the mass assignment flaw in Anything-llm's manager role creation allowed unauthorized creation of administrator accounts, scoring an 8.1, further emphasizing the need for robust access control.

SQL Injection and Data Tampering

SQL injection remains a perennial threat to API security, enabling attackers to manipulate database queries. The Cacti API demonstrated multiple SQL injection vulnerabilities, with the most severe scoring an 8.8, where attackers could inject malicious SQL commands via automation API endpoints. These vulnerabilities allow attackers to access, modify, or delete data, posing a significant risk to data integrity. The data tampering issues in NVIDIA Triton Inference Server for Linux and the SQL injection through the api_key in litellm further exemplify the critical need for secure query handling and input sanitization. These vulnerabilities underscore the necessity for prepared statements and rigorous input validation to thwart SQL injection attacks.

Improper Authorization and Session Fixation

Improper authorization mechanisms can lead to unauthorized data manipulation and access. The AnythingLLm API's vulnerability allowing managers to create administrator accounts and ZenML's session fixation issue underscore the risks associated with weak authorization controls. These vulnerabilities can lead to unauthorized access, data breaches, and potential system manipulation. In AnythingLLm, attackers could exploit improper authorization checks in API endpoints to escalate privileges, score a 7.2, and potentially take control of the system. ZenML's session fixation vulnerability, scoring a 4.2, allowed attackers to fixate sessions and bypass authentication, compromising user accounts and data integrity.

Denial of Service (DoS)

Denial of Service (DoS) vulnerabilities disrupt service availability and can be leveraged for further exploitation. The Frigate NVR's susceptibility to DoS via long Unicode filenames, scoring a 6.8, highlighted how attackers could crash the system by overwhelming it with malformed requests. This vulnerability can degrade system performance, disrupt service availability, and provide a foothold for more advanced attacks. In the NVIDIA Triton Inference Server, DoS vulnerabilities could be exploited to overload the server, causing significant service interruptions and potential data loss.

Server-Side Request Forgery (SSRF)

Server-Side Request Forgery (SSRF) vulnerabilities enable attackers to manipulate server-side requests, leading to potential internal network exploitation. The SSRF vulnerability in Lobe-chat’s API proxy endpoint, scoring a 9.0, allowed attackers to initiate unauthorized requests from the server, potentially accessing internal services and data. This type of vulnerability can be leveraged to escalate attacks, gain unauthorized access to internal networks, and exploit other internal services. SSRF vulnerabilities highlight the importance of validating and sanitizing user inputs to prevent unauthorized server-side request manipulations.

Similarities with AppSec Threats

Do these threats for AI APIs look familiar to you? The same is true for us! These vulnerabilities have been known for years as OWASP Top-10 or AppSec threats. The similarities between AI API vulnerabilities and traditional application security (AppSec) threats are striking. For instance, remote code execution, SQL injection, and improper authorization are well-known issues in the AppSec domain. These threats are not new; they have been plaguing web applications for years and now, they are resurfacing in the realm of AI APIs.

Consider SQL injection, a classic example. In both traditional applications and AI APIs, this vulnerability allows attackers to execute arbitrary SQL commands, compromising data integrity and confidentiality. Similarly, remote code execution, which can lead to complete system control by malicious actors, is a critical threat in both domains. Improper access control and path traversal vulnerabilities are also common, allowing unauthorized data access and system manipulation.

These parallels suggest that AI API security is not an entirely new frontier but an extension of established AppSec principles. The same fundamental weaknesses are being exploited, albeit in a new context. This realization has significant implications for how we approach AI API security. Instead of treating AI API vulnerabilities as a distinct category, it would be more effective to unify the OWASP Top-10 across traditional applications, APIs, and AI systems. This unified approach can streamline security practices and reduce confusion among security experts.

For example, improper authorization is a critical issue across all three areas. Whether in web applications, APIs, or AI systems, insufficient permission checks can lead to unauthorized data access and system control. By recognizing this as a single, unified threat, security teams can apply consistent mitigation strategies, enhancing overall security posture. Similarly, input validation and sanitization are fundamental practices that can prevent a wide range of vulnerabilities, from SQL injection to SSRF, across all platforms.

By unifying the OWASP Top-10, we acknowledge that these security principles are universally applicable, regardless of the specific context. This holistic approach not only simplifies the security landscape but also ensures that best practices are consistently applied. Security experts can focus on mastering a single set of guidelines, rather than navigating the complexities of separate standards for different domains.

What Surprised Us This Quarter

Surprise #1: A Surge in AI API Exploits

This quarter, we were astounded by the sheer volume of AI API exploits discovered. The number of vulnerabilities tripled compared to the previous quarter, underscoring the growing importance of securing AI systems. From remote code execution to path traversal and mass assignment, the variety and severity of these exploits highlight a critical area of concern. Notably, the Anything-llm API was particularly vulnerable, with issues ranging from arbitrary file deletion (9.9) due to path traversal in the logo photo upload feature to remote code execution using environment variables (9.6). These vulnerabilities allowed attackers to manipulate files, escalate privileges, and execute arbitrary commands, compromising the integrity and security of the AI systems.

ZenML also faced significant threats, with a directory traversal vulnerability (9.9) in its /api/v1/steps endpoint, allowing unauthorized access to sensitive files. This level of exploit suggests a rapid evolution of both AI technologies and the threats targeting them. The pace at which these vulnerabilities are being identified reflects the increasing reliance on AI-driven solutions and the corresponding rise in attack vectors as malicious actors seek to exploit these systems.

Surprise #2: So Many Issues in Grafana

Another unexpected revelation was the high number of vulnerabilities in Grafana, a widely used open-source platform known for its strong security focus and active community. Grafana's vulnerabilities this quarter were particularly surprising given its reputation. For instance, a vulnerability allowing users outside an organization to delete a snapshot with its key scored a 6.5, indicating a significant risk. Additionally, the directory traversal flaw for .csv files and multiple OAuth-related issues, including account takeover and token leakage, were alarming. These vulnerabilities highlight gaps in even the most security-conscious projects, underscoring that no platform is immune to security flaws.

Grafana's account takeover via OAuth vulnerability scored a 7.5, revealing how attackers could exploit improperly managed OAuth tokens to gain unauthorized access. The data source and plugin proxy endpoints leaking authentication tokens (7.5) further exemplified the risks associated with insufficient token management. These issues in such a popular and actively maintained project like Grafana serve as a stark reminder of the ongoing challenges in securing complex software systems.

Surprise #3: Misuse of JWT in Everywhere

Equally surprising was the continued misuse of JSON Web Tokens (JWT) across various applications, from enterprise solutions like Veeam to open-source frameworks. JWT is a widely adopted standard for securing API communications, but its proper implementation remains challenging. For instance, Veeam Recovery Orchestrator's use of a hard-coded JWT secret scored a critical 9, highlighting a fundamental security lapse. Other issues included Lua-Resty-JWT's authentication bypass (7.3), Python-jose's JWT bomb attack (6.7), and Openshift/telemeter's bypassable issuer check during JWT authentication (7.5).

In Veeam's case, the use of a hard-coded JWT secret exposed a significant vulnerability, allowing attackers to forge tokens and gain unauthorized access. Similarly, the JWT bomb attack in Python-jose exploited the decode function to overwhelm the system, potentially leading to denial-of-service conditions. These vulnerabilities indicate that developers still struggle with implementing JWT securely. The inherent complexity of JWT, combined with the need for meticulous configuration, often leads to security oversights that can have severe consequences.

Final Words: API ThreatStats Q2’24

This quarter, we decided to split our conclusions into two parts: one for CISOs and another for Security Practitioners. This approach stems from valuable feedback received on our previous reports, highlighting the need for tailored insights that address the distinct responsibilities and challenges faced by these groups. By providing focused information, we aim to better equip both strategic leaders and technical experts in addressing the evolving threat landscape. We invite you to review both sections for a comprehensive understanding of the findings and their implications.

For CISOs and Cyber Executives

1. Q2-2024 breaches highlight the need for robust API security, especially during **mergers and acquisitions**. Notable breaches include those affecting TestRail (Atlassian), HelloSign (Dropbox), Duo (Cisco), and Authy (Twilio), emphasizing the vulnerabilities introduced through acquisitions. Ensuring thorough security assessments during M&A, prioritizing API security measures, and maintaining proactive incident response strategies are essential steps to protect sensitive data and prevent future breaches.
2. Investing in API security is not only crucial for protecting sensitive data but also **makes economic sense.** Mapping API security investments to the API economy shows the potential for high returns, given the margins on API calls and transactions. A robust API security framework can safeguard these transactions, ensuring the integrity and profitability of your digital business operations.
3. The breaches at Optus and Dell illustrate the consequences of poor API security practices, such as inadequate access controls and lack of rate limiting. These incidents should serve as a wake-up call to review and **enhance API security protocols.** Similarly, the continued misuse of JWT across various platforms, including enterprise solutions like Veeam, highlights the need for thorough security training and adherence to best practices.

For Security Practitioners

1. In AI APIs, vulnerabilities such as the Anything-llm's arbitrary file deletion (9.9) and ZenML's directory traversal (9.9) showcase how attackers can exploit weak points to gain unauthorized access and control. These examples highlight the need for **rigorous input validation and robust access controls**. Similarly, the high-severity flaws in Grafana, despite its active community and security focus, emphasize that no platform is immune to security issues. Practitioners must remain vigilant and proactive in identifying and addressing vulnerabilities.
2. One key insight for practitioners is the relevance and comprehensiveness of the **OWASP Top 10 for AppSec in covering API and AI security.** The OWASP Top 10 addresses common vulnerabilities and best practices that apply universally across traditional applications, APIs, and AI systems. Relying on other specialized top 10 lists can introduce confusion and fragmented approaches to security. By adhering to the well-established OWASP guidelines, security teams can implement consistent and effective security measures across all platforms.
3. **The misuse of JWT** in various applications, from Lua-Resty-JWT's authentication bypass (7.3) to Veeam's hard-coded JWT secret (9), underscores the complexities of implementing secure token-based authentication. Security practitioners should focus on ensuring proper token management, including secure storage, rotation, and validation.

Final Final Words

In Q2 2024, we observed a dramatic increase in AI API exploits, with vulnerabilities tripling from the previous quarter. Significant breaches at organizations like Dell, Dropbox, and Twilio underscore the far-reaching impact of API vulnerabilities. This surge highlights the urgent need for comprehensive security strategies that encompass AI systems, APIs, and traditional applications.

At Wallarm, we are committed to in-depth tracking of all API security issues, manually checking each one among other CVEs since 2021. Our dedication ensures a comprehensive and accurate understanding of the evolving threat landscape, enabling us to provide detailed insights and effective security solutions. The findings from Q2 breaches and vulnerabilities reflect our meticulous approach to ThreatStats™ tracking and analysis.

Our commitment to delivering detailed analysis means that security practitioners can rely on us for rigorous input validation, robust access controls, and proper token management recommendations. By staying vigilant and proactive, we help protect against significant threats, such as those seen in AI APIs like Anything-llm and ZenML, and popular platforms like Grafana.

We extend our heartfelt thanks for reading our API ThreatStats™ Q2 2024 report. Your feedback is invaluable to us as we strive to refine our research and deliver even more targeted information. Please share your insights and suggestions with us at research@wallarm.com. Together, we can continue to enhance the security landscape and safeguard our digital future.

wallarm