# ENFORCEAUTH IDENTIFIES AUTHORIZATION FAILURES AS ROOT CAUSE OF ANTHROPIC'S THREE SECURITY INCIDENTS IN FIVE DAYS

*Analysis Shows Production-Deployable Rego Policies Would Have Prevented CMS Data Exposure, 500K-Line Source Code Leak, and npm Supply Chain Attack*



SAN DIEGO, CA, UNITED STATES, April 2, 2026 /EINPresswire.com/ -- EnforceAuth, Inc., the AI Security Fabric providing continuous authorization enforcement across applications, infrastructure, data, and AI workloads, today published a detailed technical analysis of three security incidents at Anthropic, Inc. that occurred between March 26 and March 31, 2026. The analysis identifies a single structural failure common to all three incidents — the absence of continuous authorization enforcement on non-human identities — and demonstrates through production-deployable Open Policy Agent (OPA) Rego policies how EnforceAuth's platform would have prevented each breach.

The three incidents — a content management system misconfiguration that exposed close to 3,000 internal assets including details of an unreleased AI model; a CI/CD pipeline error that published 500,000+ lines of proprietary Claude Code source code to a public npm registry; and a concurrent supply chain attack that injected a Remote Access Trojan into the same npm installation window — collectively represent the most significant public demonstration to date of what EnforceAuth terms the Authorization Gap: the

structural void between enterprise investment in AI safety and enterprise investment in AI security.


“Anthropic employs some of the most sophisticated AI safety researchers in the world. These incidents had nothing to do with AI safety. No guardrail failed. No model went rogue. What failed was the authorization enforcement layer on the service accounts, build pipelines, and package dependencies that build and ship AI products — identities that were operating with implicit trust and no continuous authorization enforcement. That is the Authorization Gap, and it exists at almost every organization deploying AI today.” — Mark Rogge, CEO and Founder, EnforceAuth

INCIDENT ANALYSIS AND ENFORCEMENT FINDINGS
EnforceAuth’s technical analysis, published as a full white paper including production-deployable Rego policies for each incident, identifies the following root causes and enforcement responses:

Incident 1 — CMS Misconfiguration | March 26, 2026
A content management system service account had no data classification enforcement, no endpoint scope restriction, and no behavioral baseline monitoring. EnforceAuth’s data classification policy engine would have evaluated the service account’s publish request against the classification of the assets being surfaced, matched a deny rule — draft content to public endpoint — and blocked the operation before any asset was exposed.
Regulatory exposure: EU AI Act Article 9, DORA Article 11, SEC cybersecurity disclosure rules

Incident 2 — Source Code Leak via npm | March 31, 2026 | 00:21–03:29 UTC
A CI/CD pipeline release process had no artifact classification gate preventing internal or debug artifacts from being included in public release packages. EnforceAuth’s pre-publication policy scan would have identified the 59.8 megabyte JavaScript source map file as a prohibited artifact type for a public registry destination and blocked the npm publish command before execution.
Regulatory exposure: DORA Articles 11 and 16, EU AI Act Article 15, SEC disclosure rules

Incident 3 — Supply Chain Attack via axios (concurrent with Incident 2) | March 31, 2026
Two malicious versions of the widely-used axios npm package, containing a Remote Access Trojan via a transitive dependency called plain-crypto-js, were available in the npm registry during the same window as the source code leak. EnforceAuth’s dependency identity verification would have flagged the publisher identity mismatch and unauthorized transitive dependency before installation, and its runtime behavioral enforcement would have blocked the RAT’s process spawning and unauthorized outbound connections at execution time.

Regulatory exposure: DORA Articles 11 and 19, EU AI Act Article 9, NIST SSDF

THE AUTHORIZATION GAP AND THE NON-HUMAN IDENTITY PROBLEM

The Anthropic incidents occurred against a backdrop of a rapidly expanding non-human identity attack surface. According to CrowdStrike's Global Threat Report 2025, non-human identities — service accounts, API keys, CI/CD bots, package managers, and AI agents — now outnumber human users by a ratio of approximately 82:1 in enterprise environments. Supply chain attacks increased 300% year-over-year in 2024, according to Sonatype's State of the Software Supply Chain report. The IBM Cost of a Data Breach Report 2024 places the average breach cost at $4.88 million, up 10% year-over-year.

"The AI industry has invested heavily in AI safety — alignment research, content guardrails, red-teaming. That investment is essential. But safety and security are different disciplines. Safety asks whether the AI model behaves as intended. Security asks whether the infrastructure through which the AI system operates is authorized, monitored, and enforced. You can have world-class AI safety systems and still have a catastrophic authorization gap. Polite AI is not secure AI."

— Mark Rogge, CEO and Founder, EnforceAuth

ENFORCEAUTH AVAILABILITY AND TECHNICAL RESOURCES

EnforceAuth reached general availability in February 2026 and is available with a free tier of 1 million authorization decisions per month. The platform is built on OPA-native Rego and integrates with existing identity providers, CI/CD platforms, package registries, cloud infrastructure, and AI platforms without requiring rip-and-replace of existing security tooling.

The full technical white paper — including production-deployable Rego policies for all three Anthropic incident scenarios, a CISO self-assessment checklist, regulatory mapping for DORA and the EU AI Act, competitive differentiation analysis, and an ROI framework — is available at enforceauth.com.

Mark Rogge
EnforceAuth
+1 612-868-7193
email us here
Visit us on social media:
LinkedIn

we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.