ISSUE #004 · WEAPONIZATION SPEED & IDENTITY BLAST RADIUS

The patch window closed. The registry pushed back.

LiteLLM went from GitHub Advisory Database index to honeypot exploitation in 36 hours, pnpm 11 added a 24-hour default delay before brand-new packages resolve, and the Context.ai to Vercel chain showed why OAuth grants need the same default friction as package installs.

> ATTACK

Sysdig · Michael Clark · 2026-04-27 · **CVE-2026-42208: Targeted SQL injection against LiteLLM's authentication path discovered 36 hours following vulnerability disclosure**

An attacker hits an internet-facing LiteLLM proxy with `POST /chat/completions` and a Bearer value that breaks out of the key lookup query. Because LiteLLM mixed the Bearer value into the SQL instead of parameterizing it, the attacker can probe column count, match the underlying query shape, and pull back rows from `LiteLLM_VerificationToken`, `litellm_credentials`, and `litellm_config`.

```
# the bypass in one header Authorization: Bearer sk-litellm' UNION SELECT api_key,NULL,NULL,NULL,NULL FROM "LiteLLM_VerificationToken"--
```

Timeline. Apr 20, 21:14 UTC: GHSA-r75f-5x8p-qvmc published on the LiteLLM repo. Apr 24, 16:17 UTC: indexed in the global GitHub Advisory Database. Apr 26, 04:24 UTC: Sysdig's honeypot logs first exploitation. 36 hours, 7 minutes from index to weaponization. Source IPs `65.111.27.132`, `65.111.25.67` (AS200373, 3xK Tech, Germany). User-agent `Python/3.12 aiohttp/3.9.1`, Bearer header containing `UNION SELECT`. A patch window measured in days does not survive that timeline.

▸ MONDAY CHECK

Upgrade LiteLLM to v1.83.7+ this week. If you cannot patch immediately, set `general_settings.disable_error_logs: true` (per advisory, stops the in-band leak), then put the proxy behind NGINX/Envoy and block Authorization headers containing `'`, `(`, `)`, `UNION`, `SELECT`, `FROM`, `OR`, or `--`. Rotate every virtual key, master key, and provider credential the instance ever held; grep access logs for `sk-litellm'`; audit provider billing for unexpected `/chat/completions` use.

> RULE

Socket · Sarah Gooding · 2026-05-04 · **pnpm 11 Adds Supply Chain Protection Defaults for Minimum Release Age and Exotic Subdependencies**

pnpm 11 ships two defaults that change the worm-class attack math: `minimumReleaseAge=1440` (a 24-hour delay before a newly published version resolves) and `blockExoticSubdeps=true` (transitive deps from git or tarball sources are refused unless explicitly allowed). A new `allowBuilds` model replaces the older lifecycle-script settings. The highest-risk window for a malicious package is the first few hours after publish, and pnpm 11 makes teams opt out of that delay instead of opt in. If you use npm or yarn, recreate it in your private registry or CI wrapper by blocking package versions published in the last 24 hours, and make exceptions explicit instead of silent.

> DEFENDER

Push Security · Dan Green · 2026-04-23 · **Unpacking the Vercel breach: A cautionary tale for Shadow AI and OAuth sprawl**

The chain: Context.ai was reportedly compromised, the attacker reused OAuth tokens stored in Context.ai's Supabase, and a broad `https://www.googleapis.com/auth/drive` grant on Context.ai's OAuth client (`110671459871-30f1spbu0hptbs60cb4vsmv79i7bbvqj.apps.googleusercontent.com`) exposed downstream Workspace accounts, including a Vercel employee. Open Google Admin → Security → Access and data control → API controls, review that client ID, and decide whether *"Don't allow users to access any third-party apps"* should be your default.

GitGuardian · Dwayne McDaniel · 2026-04-27 · **Short-Lived Credentials in Agentic Systems: A Practical Trade-off Guide**

Short TTL alone is not the answer (5-15 min user-facing, 15-60 background, 1-6h autonomous, all with caveats: refresh failures cause partial writes, vault availability is a single point, and the long-lived "exception" credential drifts into the default). The actual control is per-task, per-stage scoping. Same lesson as the OAuth chain on a different surface: narrower scope, shorter lifetime, explicit approval path.

> RADAR

GitGuardian · Gaetan Ferry · 2026-04-28 · **The Bot Left a Fingerprint: Detecting and Attributing LLM-Generated Passwords**

Markov-chain analysis trained on 8,000 passwords across 40 LLMs identifies AI-generated passwords with 55% model-attribution accuracy, 65% provider attribution. GitGuardian found 28,000 such passwords in a 34M-password GitHub sample (2025-11 to 2026-03), with Anthropic, Qwen, and Google the top three attributions. Same control theme as the rest of this issue: if AI tooling can mint reusable secrets, make those secrets short-lived and narrow-scoped before statistical bias turns into reusable access.

SOURCES

Sysdig / sysdig.com/blog/cve-2026-42208-targeted-sql-injection-against-litellms-authentication-path-discovered-36-hours-following-vulnerability-disclosure
Socket / socket.dev/blog/pnpm-11-adds-new-supply-chain-protection-defaults
Push / pushsecurity.com/blog/unpacking-the-vercel-breach
GitGuardian / blog.gitguardian.com/short-lived-credentials-in-agentic-systems-a-practical-trade-off-guide
GitGuardian / blog.gitguardian.com/the-bot-fingerprint-detecting-llm-passwords