Over 29 million secrets were leaked on GitHub in 2025, and AI really isn't helping

techradar.com/pro/security/over-29-million-secrets-were-leaked-on-github-in-2025-and-ai-really-isnt-helping

Sead Fadilpašić March 18, 2026

Add us as a preferred source on Google
Subscribe to our newsletter

- **GitGuardian report warns AI-driven coding leaks secrets at record pace**
- **2025 saw 29M exposed credentials on GitHub, +34% YoY jump**
- **AI-assisted commits double baseline leak rates, with MCP configs fueling exposures**

Vibe-coding may seem great for quickly shipping products, but inexperienced developers are leaving gaping cybersecurity holes that are causing breaches and exposures left and right. This is according to GitGuardian's latest report, the "State of Secrets Sprawl" paper that was just released.

In the research document, the organization said 2025 was the year when AI adoption "permanently changed" software engineering. That year, there was a 43% increase year-on-year in public commits, growing at least two times faster than before.

An increase in commits also means an increase in secrets and since 2021, these have been growing roughly 1.6 times faster than the active developer population. Also, secret leak rates in AI-assisted code were roughly double the GitHub-wide baseline.

ClaudeCode, MCP configurations, and other risks

“Together, these forces drove a +34% YoY increase in newly leaked secrets on GitHub, reaching ~29 million secrets detected overall, marking the largest single-year jump ever recorded,” the organization said in a press release.

Of all the different vulnerabilities that can be found in AI-generated code, exposed credentials remain the biggest path to compromise, GitGuardian says. Commits built with Claude Code apparently leaked secrets at roughly 3.2% which is two times the baseline, and AI service credentials leaks seem to be accelerating the fastest. Leaks tied to AI services spiked 81% year-on-year, and are “more likely” to slip through protections.

GitGuardian specifically singled out Model Context Protocol (MCP) configuration risk. The report says that MCP server documentation often recommends putting credentials in configuration files, which is a risky pattern that contributed to more than 24,000 secrets being exposed.

The paper further explained that internal repositories are six times more likely to contain hardcoded secrets, compared to public ones, and stressed that more than a quarter (28%) of incidents originate from leaks in collaboration and productivity tools.

Finally, with AI agents getting deeper local access, prompt injection and supply-chain attacks are getting more disruptive:

“AI agents need local credentials to connect across systems, turning developer laptops into a massive attack surface. We built our local scanning and identities inventory tool to protect them. Security teams need to map out exactly which machines hold which secrets, surfacing critical weaknesses like overprivileged access and exposed production keys.” said Eric Fourrier, GitGuardian CEO.