The 2026 State of Code Quality

A simple look at AI-assisted repositories

We analyzed 1,961 repository scans to uncover how AI is impacting code quality, security, and maintainability.

FIRST SCAN REALITY CHECK

38.4% of first-time scans land in D or F

Based on 1,961 unique repositories, nearly 4 in 10 codebases fail basic hygiene checks on day one.

CODE QUALITY AT A GLANCE

First-time scans are bimodal

28.3% Grade A

24.7% Grade F

Across 1,961 repositories, A's and F's are the two most common first-scan outcomes — more repos get an F than a C.

Grade C is the smallest cohort

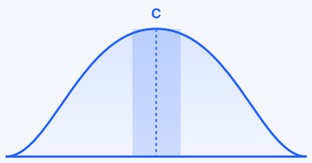

Most projects aim for strong standards (A/B) or remain in riskier territory (D/F). Only 15.9% land in Grade C.

Code quality is independent of repository age

ρ = -0.008

Analysis shows a near-zero correlation between repository age and overall score, debunking the "legacy code is worse" myth.

THE AI COHORT GAP

2.7× THE RATE

42.6% AI-Marked Repositories

15.7% Non-AI Repositories

AI repos commit secrets 2.7× more often than non-AI repos.

42.6% of AI-marked repositories contain hardcoded API keys or tokens vs. 15.7% for non-AI.

THE AI QUALITY PENALTY

AI repos score 7 to 8 points lower

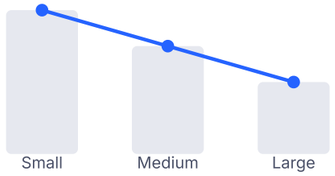

After controlling for codebase size, AI-marked repositories score lower in every size bucket.

AI-GENERATED CODE

Skews toward larger files

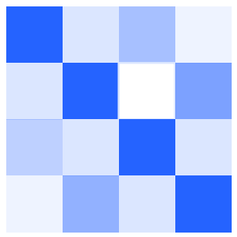

AI tools tend to append code, leading to higher line counts and increased complexity vs. manual refactoring.

GRADE DISTRIBUTION: AI VS. NON-AI REPOS

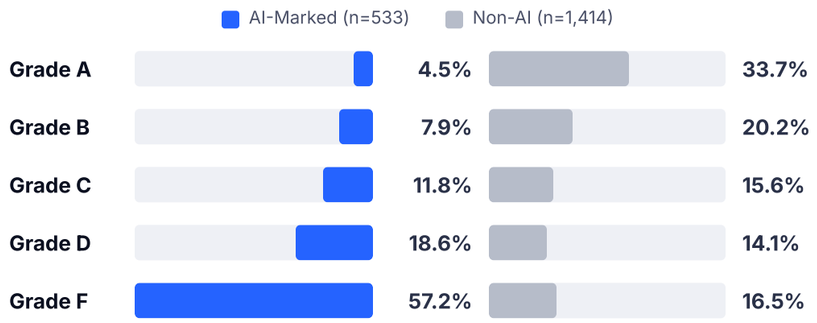

57.2% of AI-marked repos earn an F vs. 16.5% of non-AI — a 3.5× gap. At Grade A the gap flips: 33.7% vs 4.5%.

ADOPTION & RISK PATTERNS

96% of AI-marked repos fail the test-presence check

AI tools default to the shortest path that compiles — tests get skipped unless asked for explicitly.

JavaScript leads AI adoption at 46.6%

JS developers use AI tools at 3.8× the rate of Rust developers, reflecting a high-velocity "MVP" culture.

JS

9 in 10 scans finish in under 17 seconds

99 in 100 finish in under 75 seconds. Code-quality tooling isn't a budget-blower anymore — it's faster than a git clone.

AI COHORT RISK INTENSITY

(Finding-count ratio vs. non-AI baseline)

3.20×	2.84×	2.81×	2.33×
Hardcoded Secrets	Complexity Issues	Error-Handling Gaps	Missing Tests

The AI cohort fires every check at 2–3× the non-AI baseline across the full 1,961-repo sample.