

Review Debt and What It Costs

2026-07-12 / 00:12:27

"These metrics tell you the speed of production, not the speed of trust." — Sachin Gupta

— Seln Oriax, today's narration

Sachin Gupta at eBay just released the first quantifiable framework for measuring review debt — the compounding gap between agent-produced code and code humans have actually reviewed.

Roberto Stagi reports TypeScript has surpassed Python as GitHub's most popular language, driven by coding agents defaulting to JS. Terry Tao used an agent to port 24 Java applets from 1999 and found only one bug — a bounded task done well. Starbucks is rebuilding \$400 million in Microsoft and IBM software in-house.

All of it points to the same tension: agents ship code faster than trust can keep up, and nobody's measuring the interest rate.

CHAPTERS

00:00:04 The Gap Nobody's Measuring

00:03:03 The Language That Agents Default To

00:05:18 Bounded Work, Bounded Risk

00:07:33 Where the Boundary Should Be

00:09:18 The Buy-vs-Build Ticking Clock

00:11:23 The Thread Between Them