WHAT UBER'S CTO COULDN'T ANSWER

The AI ROI Gap

Adoption is easy to encourage. Return is almost impossible to prove.

THE DATAPOINT

Uber's full-year 2026 AI budget was exhausted by April.

"I'm back to the drawing board, because the budget I thought I would need is blown away already."

Praveen Neppalli Naga, CTO, Uber. The Information, April 2026.

84%
of Uber developers classified as agentic coding users by March 2026.

6x
increase in AI-related costs at Uber since 2024.

11%
of live backend code updates now written by AI systems.

Why this isn't an Uber problem. It's an enterprise problem.

Every CFO funding AI tooling in 2026 is staring at the same gap. Adoption telemetry lives in vendor consoles. Productivity output lives in Git, Jira, Linear, and Salesforce. Cost lives in finance. No single system reconciles them. That is why a $3.4B R&D org with dedicated planning teams missed its forecast inside one quarter.

Internal leaderboards reward usage. Usage is not the same as contribution. Without a layer that ties consumption to delivered output and quality, AI spend compounds invisibly until a budget cycle exposes it. By then it is too late to govern.

Three questions a CTO and CFO cannot answer today

- **Adoption depth.** Who is actually using AI tooling, and where is spend concentrated?
- **Productivity contribution.** Has cycle time, throughput, or feature velocity actually moved?
- **Cost per outcome.** What does it cost to ship a PR, close a ticket, or deliver a feature, by team?

There is a layer missing between vendor consoles and the CFO's spreadsheet.
Levos closes that gap. See the snapshot on the next page.

SAMPLE OUTPUT / ENGINEERING ORG, 5,000 SEATS

AI Impact ROI Snapshot

Five blocks. Every number sourced. Every number drillable to underlying signal.

1 ADOPTION DEPTH

Who is using AI tooling, and where is spend concentrated.

ACTIVE USERS	POWER USERS (TOP 10%)	DORMANT LICENSES	AVG TOKENS / DEV / WK
4,213	**421**	**787**	**1.42M**
of 5,000 seats | 84%	consume 62% of tokens	no use, last 30 days | 16%	+38% QoQ

2 PRODUCTIVITY LIFT

Throughput delta vs pre-AI baseline. Cross-referenced with Delivery family.

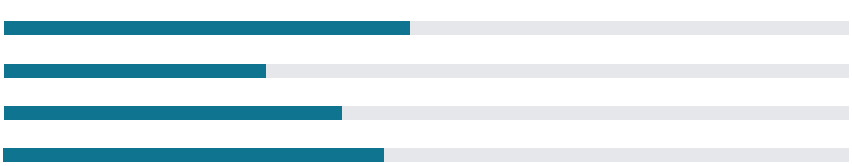

3 QUALITY TRACE

Defect, revert, and escape rates on AI-assisted vs human-only PRs.

REVERT RATE (AI-ASSISTED)	DEFECT ESCAPE RATE	TEST COVERAGE DELTA
4.8%	**+12%**	**-2.1pp**
vs 3.1% human	QoQ on AI PRs	in AI-heavy services

4 COST PER OUTCOME

AI spend per delivered unit, by team.

TEAM	$ / PR	$ / TICKET
Platform	$148	$62
Payments	$214	$94
Growth	$382	$176
Mobile	$96	$41

5 CONCENTRATION & RISK

High spend, low Quality teams. The kill list.

Growth	High spend, declining quality
Data Eng	High spend, flat throughput
Mobile	Low spend, strong lift
Platform	Balanced

Want this snapshot built on your data? Pilot a Levos AI Impact assessment.
30-day pilot, your existing connectors. Contact: hello@levos.info | levos.info