

When AI Writes Most of the Code: The C-Suite Executives Briefing



Breaking news and context. Anthropic reports that most of the code for recent Claude releases is AI-authored, with engineers acting as reviewers, integrators, and control owners. **A fundamental shift in how software is built, and how it must be governed.**

AI can now draft a large share of software changes. Your upside is faster delivery at lower unit cost. Your risk is defects, drift, and spend blow-outs if governance is weak. The fix is simple: treat AI-assisted coding as an operating change, not a developer toy.

I am writing this as a former Head of PMO, Head of Operations, and transformation lead. **My lens is delivery discipline: how to convert this shift into speed with safety, audit, and measurable ROI.**

What you get (in business terms)

- **Speed.** Shorter lead times from idea to release. Backlogs shrink. Months become weeks.
- **Cost per change.** Less manual coding and rework. Fewer hand-offs. Lower unit cost.
- **Consistency.** Fewer “human style” differences. Easier to maintain and audit.
- **Focus.** Engineers spend time on design, security, and integration, not boilerplate.

Where the money shows up

- **Run-rate savings.** Same team, more changes shipped. Or same output with a smaller bench.
- **Revenue timing.** Features that move a KPI arrive earlier. Cash comes forward.
- **Risk and quality.** Fewer production incidents and faster fixes lower warranty and reputational cost.
- **Change velocity.** Faster experiments mean faster “stop or scale” decisions.

What can go wrong (and how to prevent it)

- **Hidden rework.** Fast generation creates slow clean-up later.
 - **Countermeasure:** write the test before the change; do not merge code that fails it.
- **Security and data leaks.**
 - **Countermeasure:** set “least-privilege” access; log every AI action; ban write actions unless tests pass.
- **Cost drift.** Token/CPU usage creeps up quietly.
 - **Countermeasure:** set a cost ceiling per change; fail releases that exceed it.
- **Irreversible rollouts.**
 - **Countermeasure:** switch everything on with a **feature flag**; roll out to 1–5% first (**canary**); auto-rollback on threshold breach.

The governance shift (non-technical)

- **One KPI per initiative.** Example: “Cut average release lead time from 15 to 9 days by 30 June.”
- **One owner.** Name the accountable person. No committees.
- **Go/Kill rule.** If the KPI is not moving inside the agreed cost/risk limits within 60 days, stop or re-scope.
- **Audit trail.** Record which model produced what code, who reviewed it, and which tests passed.
- **Dashboards you can read.** Lead time, escaped defects, exceptions fixed within SLA, cost per change.

A simple 90-day plan

Days 0 - 30 - Pick one thin slice

- Choose a safe, high-volume change type (reports, internal APIs).
- Set a KPI, cost ceiling, and owner.
- Agree success/failure criteria in writing.

Days 31 - 60 - Prove it on real work

- Use AI to draft changes; reviewers approve against pre-written tests.

- Run **shadow** mode first (no user impact), then a **canary release** (1–5% of traffic).
- Publish “before vs after” on speed, defects, and cost per change.

Days 61 - 90 - Scale by evidence

- Expand to a second slice.
- Add model/vendor swap-ability, so you are not locked in to one provider only.
- Put the metrics on a monthly COO/CFO pack.

Back-of-the-envelope economics (illustrative)

- If a team of 8 engineers costs ~£100k/month.
- Today they ship ~40 production changes/month. Unit cost ≈ **£2,500/change**.
- With AI assist and tests-first, they ship 60 changes/month at similar quality.
 - Unit cost drops to **~£1,670/change** (–33%).
 - Tool + compute adds **£6 - 10k/month**. Net saving ≈ **£24 - 28k/month, £288 - 336k/year**.
- If incidents fall 20%, add avoided outage and reputational value on top.

Numbers vary. **The method is what matters: measure unit cost, quality, and speed together.**

What to ask your team (board-level questions)

1. Which one KPI are we moving, by when, and how will we stop if it does not move?
2. Where is AI drafting code today, formally or informally? Who signs off?
3. Show me a change with its audit trail: tests, reviewer, cost, and rollback plan.
4. What is our cost per change trend since introducing AI assist?
5. Can we switch model or vendor without rebuilding the product? When was that last tested?

Procurement and policy in plain terms

- **Buy governance.** Demand features that support audit trails, cost limits, and easy rollback.
- **Avoid lock-in.** Require a model-agnostic layer so you can change supplier without rewiring your product.
- **Data policy.** Keep sensitive data local where possible; be explicit about what leaves and why.

Where to start (examples)

- **Finance/ops:** nightly file normalisation. AI drafts the mapping; tests must pass before merge.
- **Customer ops:** templated responses with live data pulls; agent proposes, human approves.
- **Compliance:** document classification and field extraction; exceptions only to a human.

Bottom line

This is not about “more AI”, but about **faster, cheaper, safer transformation**.

Write the test first. Let AI draft. Review against the test. Roll out small. Roll back fast. Measure unit cost, quality, and speed on one page.

Contact me if you want a 30-minute run-through, I can create a one-page plan and a thin-slice candidate you can ship in 6 - 8 weeks-no hype or theatre, just outcomes.

#AI #DigitalTransformation #PMO #OperationalExcellence #Governance #RiskManagement

This article draws on analysis of a range of articles and news sources relating to AI and technology.

Sources

Anthropic CEO Said AI Would Write 90% of Code in 3-6 Months <https://www.remio.ai/post/anthropic-ceo-said-ai-would-write-90-of-code-in-3-6-months-here-s-what-the-evidence-shows>

Dario Amodei: AI Will Write Nearly All Code in 12 Months!!
https://www.reddit.com/r/OpenAI/comments/1j8r0jw/dario_amodei_ai_will_write_nearly_all_code_in_12/

Claude Code Coded Claude Cowork - The Algorithmic Bridge <https://www.thealgorithmicbridge.com/p/claude-code-coded-claude-cowork>

AI to write 90% of code in 3-6 months? | Alex Banks posted on the ... https://www.linkedin.com/posts/a-banks_anthropic-ceos-shocking-prediction-90-activity-7306298173008805888-Es_d

Most Anthropic Teams Are Coding With AI, but It's Not Replacing ... <https://www.businessinsider.com/most-anthropic-teams-coding-with-claude-ai-not-replacing-humans-2025-10>

How 2026 Could Decide the Future of Artificial Intelligence <https://www.cfr.org/article/how-2026-could-decide-future-artificial-intelligence>

When AI writes almost all code, what happens to software ... <https://newsletter.pragmaticengineer.com/p/when-ai-writes-almost-all-code-what>

Anthropic Shipped Cowork in 10 Days Using Its Own AI. Here's Why ...
<https://karozieminski.substack.com/p/claude-cowork-anthropic-product-deep-dive>

Anthropic Takes Claude Code to the Web | The Tech Buzz <https://www.techbuzz.ai/articles/anthropic-takes-claude-code-to-the-web>

Anthropic opens Claude Code's automation power to everyone with ... <https://ppc.land/anthropic-opens-claude-codes-automation-power-to-everyone-with-cowork/>

90% of Claude's Code is Now AI-Written: Anthropic CEO Confirms ...
<https://markets.financialcontent.com/wral/article/tokenring-2026-1-13-90-of-claudes-code-is-now-ai-written-anthropic-ceo-confirms-historic-shift-in-software-development>

Anthropic CEO Predicts AI Code Writing by March 10, 2026 - LinkedIn
https://www.linkedin.com/posts/stephenpontes_its-officially-2-months-march-10-2026-activity-7412927769073864706-jar3

Is 90% of code at Anthropic being written by AIs? - LessWrong
<https://www.lesswrong.com/posts/prSnGGAgfWtZexYLp/is-90-of-code-at-anthropic-being-written-by-ais>

AI coding is now everywhere. But not everyone is convinced.
<https://www.technologyreview.com/2025/12/15/1128352/rise-of-ai-coding-developers-2026/>

AI Writes 100% of My Code and 0% of My Code - Future History <https://danieljeffries.substack.com/p/ai-writes-100-of-my-code-and-0-of>

Disclaimer and Disclosure

Third-party Content and AI Assistance: This article references tools and software that are publicly available and proprietary to their respective creators. The author does not claim ownership or affiliation with these third-party products. This article was written by the author with assistance from Generative AI Language Models.

Transparency Notice: While every effort has been made to ensure accuracy, readers should verify information independently and consult official sources or documentation for the mentioned tools and software. The use of AI in the writing process is disclosed in the interest of transparency, but all opinions and analyses are the author's own unless otherwise stated.