

Hands up!

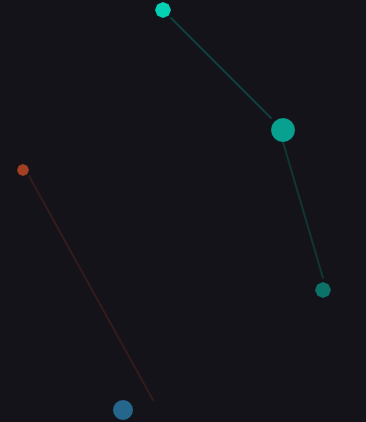
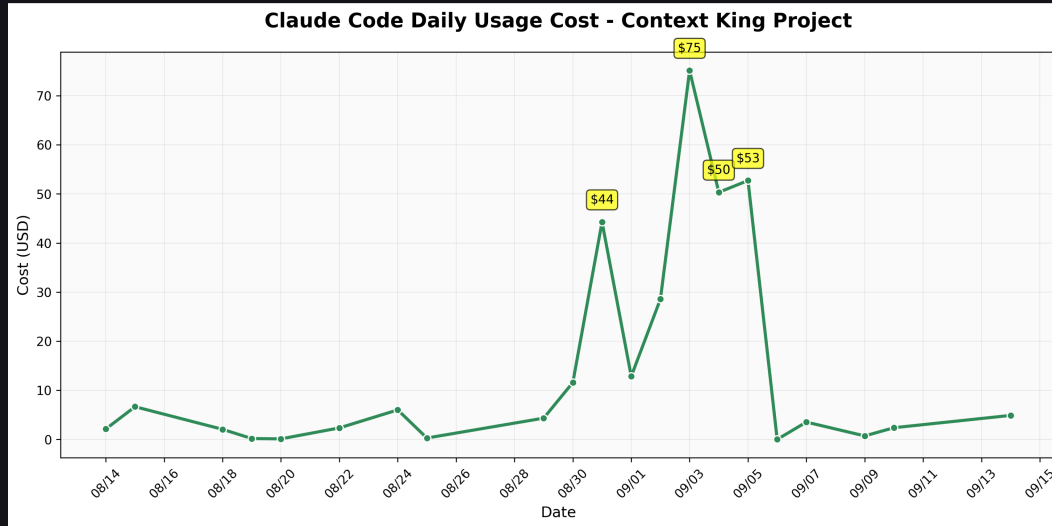
You've used a coding assistant?

You have an active AI subscription?

You use AI daily?

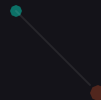


⚠ True Story



ABOUT ME

Bauke Brenninkmeijer



Applied Researcher @ [orq.ai](#)



Data Science + Computer Science Background

ABN AMRO, ING



Organise MLOps Community Amsterdam

AI focused meetup (2.5k members)



Build AI + Uses AI to Build + Evaluate AI

Wrong about this from several sides

ABOUT ME

How do i use AI

Whole company has Claude Code Max

Many greenfield projects

All PRs and bulk of reviews are AI assisted

Recurring tasks are automated as much as possible

The Trust Gap

Why hesitation comes from unfamiliarity

50%

Trust increase in 2 years

26-46pp

Higher trust with hands-on



The Trust Gap

Why hesitation comes from unfamiliarity

50%

Trust increase in 2 years

26-46pp

Higher trust with hands-on



Who has engineers using AI tools without telling you?

"MY COMPANY WON'T LET US"

Making the Case to Leadership

90%

of engineering teams use AI tools

78%

of orgs use AI in at least one function

1 in 3

work around company AI policies

Compliance headache - but it's a signal your tooling and policies haven't kept up with demand.

Who has had requests for tools blocked by policy?

Connection to DORA Metrics

DORA METRIC	HOW AI HELPS	HOW AI HURTS (without guardrails)
Deploy Frequency	Faster feature delivery, more PRs shipped	More half-baked PRs getting merged
Lead Time	Faster first drafts of code and tests	PR review time up 91% (Faros AI)
Change Failure	AI reviews catch issues pre-deployment	Confident-but-wrong code ships unreviewed
Time to Restore	AI-assisted triage + fix suggestions	Over-reliance on suggestions wrong 47% of time

21% more tasks
Individual developer gains (Faros AI)

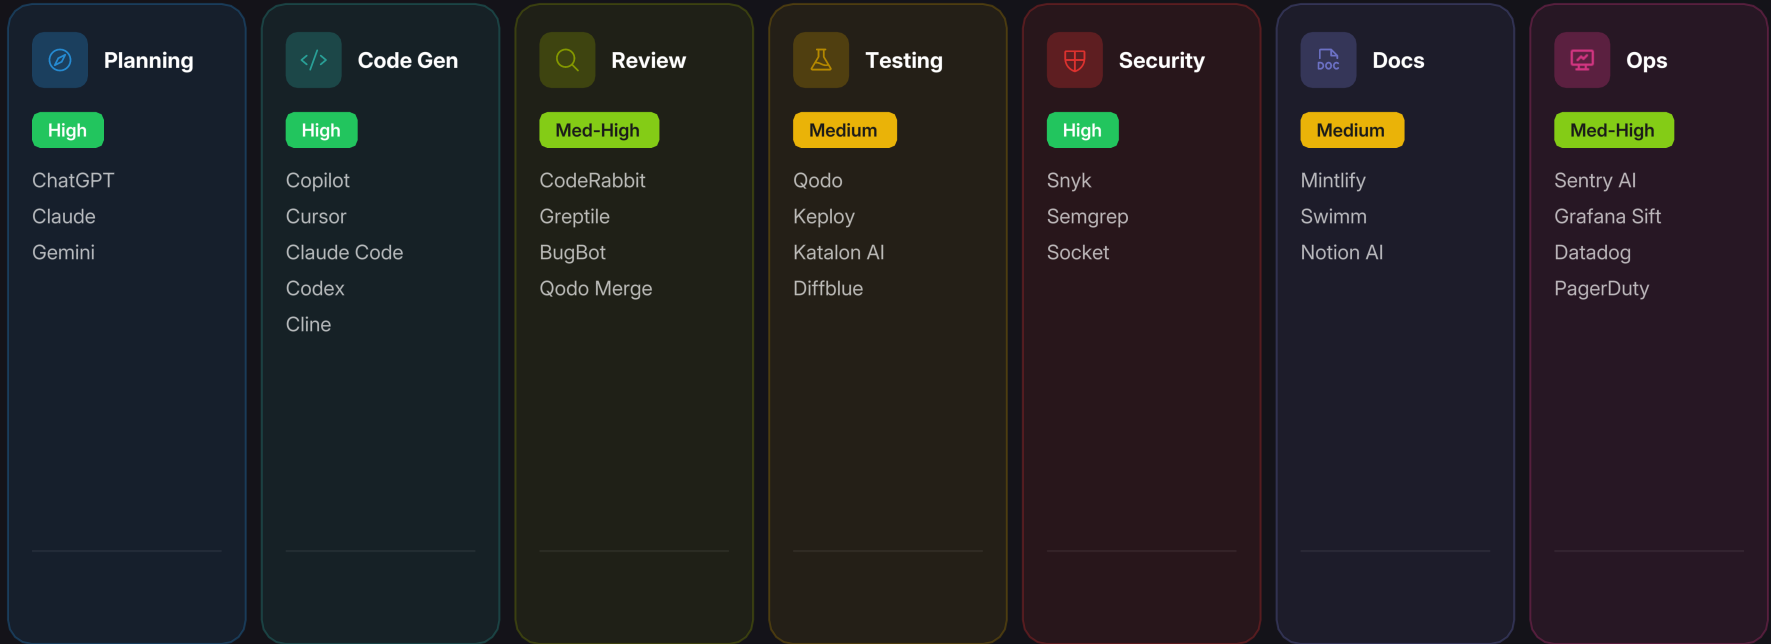
98% more PRs

But org-level delivery metrics stayed flat.

Why? Coding is only 16% of dev time. Even big speedups are bounded by review, coordination, and understanding requirements.

WHAT'S OUT THERE

AI Tools Mapped to the SDLC



This map is constantly changing. Tools are added weekly. Anthropic is aggressively entering almost all stages of the SDLC.

AI application ontology



Vertical Vendor Solutions

CodeRabbit
Greptile
BugBot
Qodo Merge

Internally facing AI applications

Standup automation
Sales prospect exploration
Invoice automation
KYC

Coding Assistant

Shadow IT

Externally facing AI applications

Customer chat
Personalized marketing
Customer recommendations
Customer lifecycle management

THE PARADIGM SHIFT

The Coding Assistant Is the Hub

Not separate UIs — one interface connects engineers to the entire AI toolchain



MCP (Model Context Protocol) is accelerating this convergence — tools plug into the assistant.

FROM THE TRENCHES

The 80/20 Problem

80%

fast



20%

the real work



KEY INSIGHT

This is where most people get frustrated and give up.

The skill: knowing when to stop prompting and start editing.

Like Google Maps: great for the neighborhood, but for the last 50 meters you're better off just looking around.

Who recognizes this pattern?

What Gets Faster vs What Breaks

What Gets Faster

- ✓ Boilerplate, scaffolding, repetitive patterns
- ✓ Exploring unfamiliar codebases
- ✓ Writing tests for existing code
- ✓ First drafts of documentation
- ✓ Greenfield projects (less context)

What Breaks

- ✗ Confident but wrong assumptions
- ✗ Complex existing codebases
- ✗ Nuanced business logic
- ✗ Delete tests that fail
- ✗ Tests that pass but don't verify

66%

Say solutions "almost right"

"Review matters more with AI, not less."

How to make it work?

WHAT DIDN'T WORK

- ⊗ Manual review only - too slow, too much
- ⊗ Blanket ban - drove usage underground



Input/Output Filtering

Control what goes in and comes out of the model. Block sensitive data, enforce patterns.



Eval Frameworks

Measure quality systematically, not anecdotally. Run evals on every PR.



Human-in-the-Loop

Define where in the workflow humans verify. Not everywhere - just the critical checkpoints.

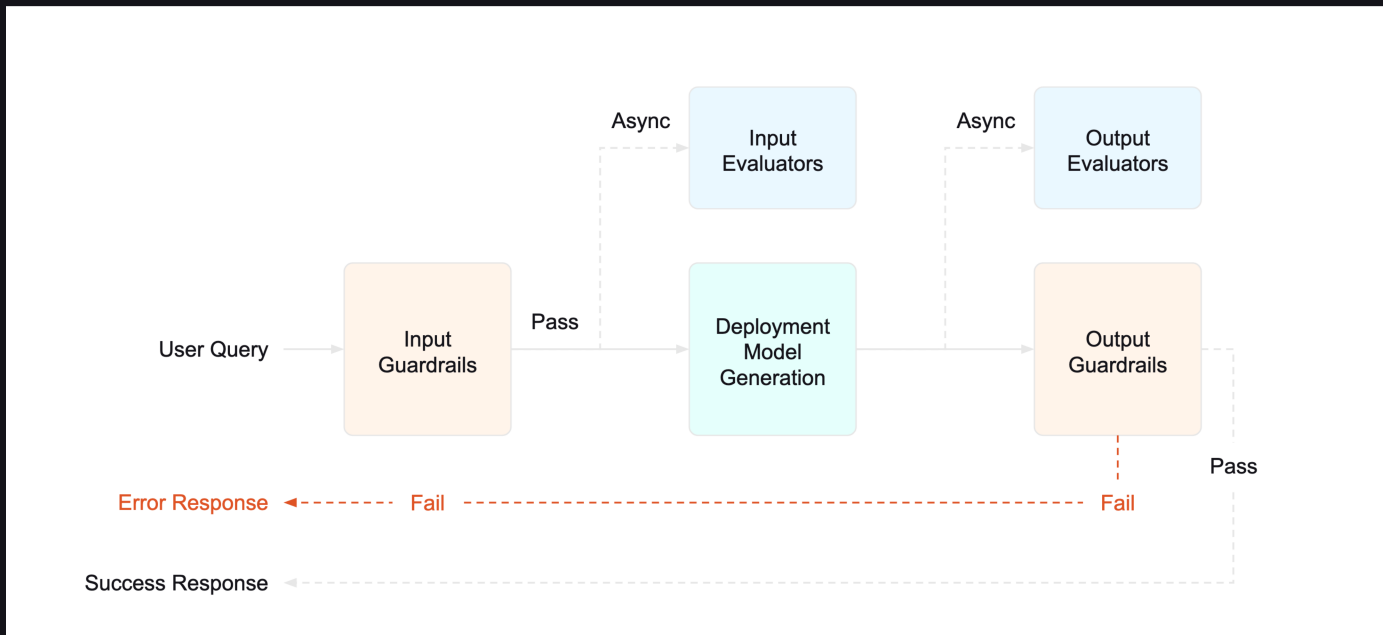


Audit Trails

Who generated what, when, and what was changed? Visibility creates accountability.

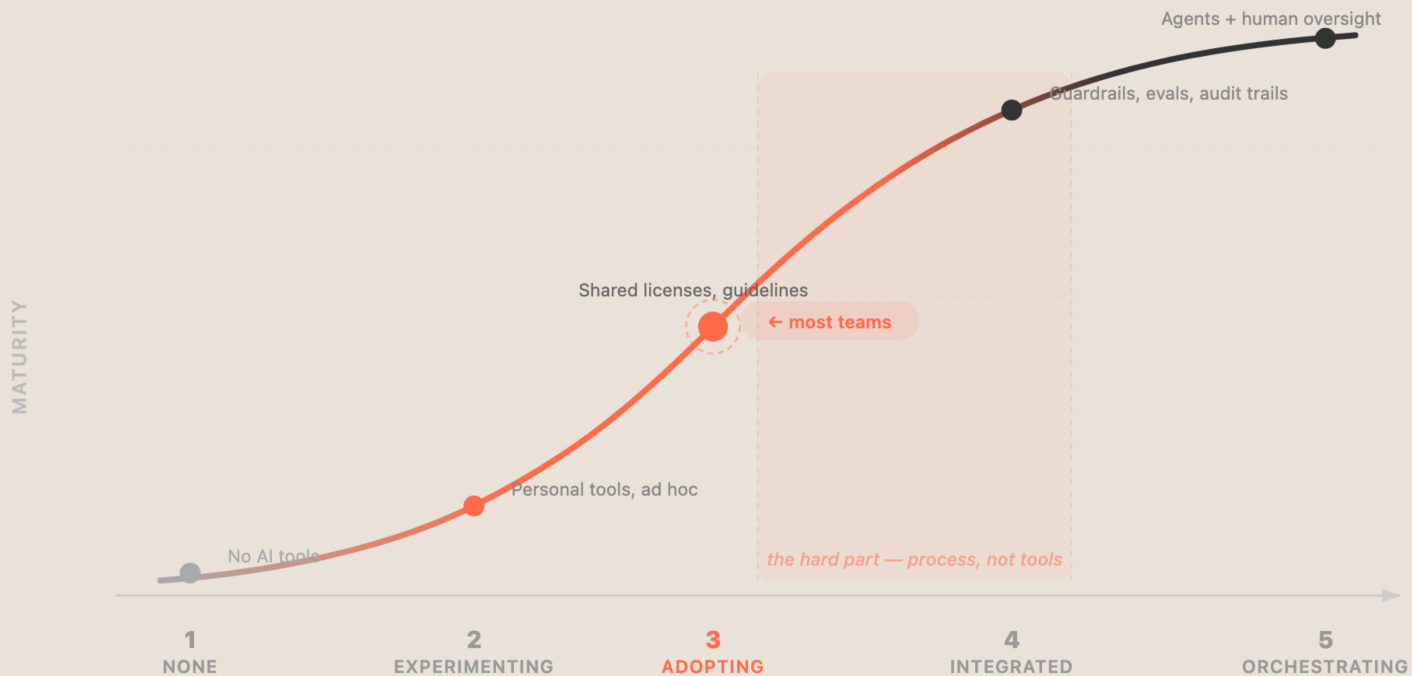
The point isn't our specific tools - it's that this can be systematized.

What Actually Works






The point isn't our specific tools - it's that this can be systematized.

AI Maturity Curve



AI Maturity: Score Yourself

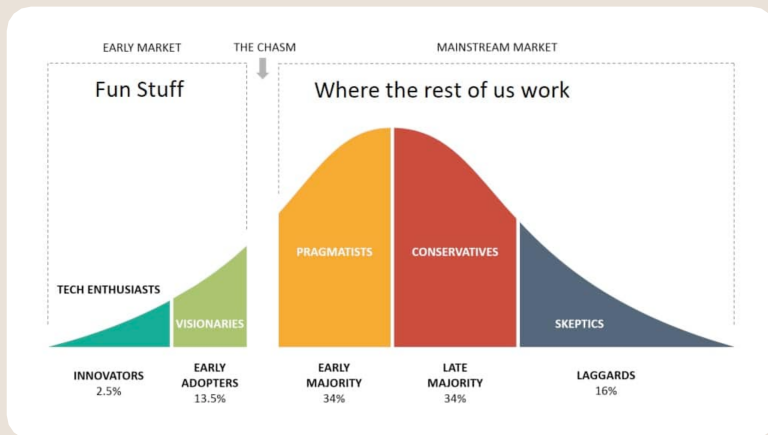
Score each dimension 1-5. Be honest - this is for you, not for a report.

 Tool Adoption	1 No AI tools	2 Personal use	3 Team licenses	4 In CI/CD	5 Agents
 Governance	1 No policy	2 Shadow AI	3 Guidelines	4 Guardrails + evals	5 Measured + enforced
 Team Fluency	1 No exposure	2 Few enthusiasts	3 Most trying	4 Know when (not) to	5 Directing agents

Add your scores, divide by 3. That's your level.

Now: where's the biggest gap between your dimensions?

Getting Your Team On Board



15-20%

Enthusiasts

Already using it,
sometimes secretly

60-70%

Pragmatists

Open to it if shown clear value

15-20%

Skeptics

"Could have written
it faster myself"

WHAT WORKS

- Pair programming sessions with AI tools
- Dedicated experimentation time
- Share wins visibly (demos, Slack)

WHAT DOESN'T WORK

- Top-down mandates without support
- Measuring lines of AI code
- Ignoring security concerns

Focus on the pragmatists. The enthusiasts will help you.

Taking the First Steps

Start in shadow mode. Build trust. Then scale.

SHADOW MODE ROLLOUT

1 Shadow

AI runs alongside humans. Every output is verified. Nobody relies on it yet - just observing.



2 Assist

AI suggestions shown to humans. They accept, edit, or reject. Trust builds through experience.



3 Automate

Proven tasks run autonomously. Humans review exceptions only.

START WITH BORING WINS — LOW RISK, HIGH VISIBILITY



Standup summaries

AI summarizes daily Slack updates



Ticket triage

Auto-categorize support requests



Test generation

Generate tests for legacy code



PR descriptions

Auto-generate from commit diffs

The goal: build enough trust that "let AI handle this" becomes a natural team instinct.

The picture is more nuanced than "AI = faster"

THE CONCERNS

32.7% AI PR acceptance rate
vs 84.4% human PRs — LinearB, 8.1M PRs

-19% Slower for new AI users
METR study: devs with under 50 hours experience

60% Positive sentiment (was 70%+)
Stack Overflow 2025 — honeymoon is over

THE PROMISE

+21% Faster task completion
Google RCT - 96 Engineers

2x PRs for daily AI users
DX Q4 2025, 135K devs — vs light users

4.4h Saved per week (Staff+ engineers)
Experienced devs gain the most from AI tools

Models keep improving

New releases every few months. Your experience compounds — the same skills work on better models, multiplying your leverage over time.

💡 Teams tackle harder problems

NAV IT and Agile Teams studies: output stayed flat, but teams shifted to higher-complexity work. The value is qualitative, not just speed.

Where This Is Going

THE SHIFT

"AI deprecates skills like language expertise, but amplifies vision, strategy, task breakdown, and feedback loops."

— Kent Beck

THE SKILLS THAT MATTER NOW



Task Decomposition

Breaking work into pieces an AI agent can handle independently



Context Engineering

Giving AI the right constraints and success criteria upfront



Review and Verification

Catching confident-but-wrong output before it ships



Product Thinking

Understanding what to build matters more when how gets cheaper

How Will Your Role Be Redefined?



Have you thought about how AI changes what you do daily?

Not just the tools — the actual shape of your work



Raise your hand

Have you updated the interview process to reflect AI-assisted work?



Automated tasks

Which parts of your workflow has AI already taken over?



New skills

What skills are you developing that didn't exist 2 years ago?

The role of engineer is being redefined in real-time. Are you shaping it, or reacting to it?

Thank You

Bauke Brenninkmeijer

Applied Researcher @ orq.ai

Questions? Let's chat after.



AI in Your SDLC | iO Engineering Manager Community

