

The Unsexy Truth About AI Success

Leadership, Hygiene, and Human Foundations





John Harbison

Head of Product Engineering at
PrescriberPoint, pioneering
agentic workflows in healthcare.

<https://www.linkedin.com/in/john-harbison/>



Jess Wolfe

Customer Success Manager at
Swarmia, focused on driving
engineering effectiveness.

<https://www.linkedin.com/in/thejessicawolfe/>



The question that started this...

**How do you measure
engineering effectiveness now
that AI is in the picture?**

...the answer is not story points.



AI amplifies
what you
already have.



You cannot
outrun process
breakdowns by
adding more
AI.



The Measurement Baseline (Camry vs. IndyCar)



Metrics designed to keep the driver calm. Lots of problems are hidden until a threshold triggers a warning.



Metrics designed to protect a highly stressed engine and maximize performance. Live vitals and alarms are front-and-center, plus pit wall telemetry.



Measurement Baseline: Traditional vs AI-Optimized

Camry (Traditional)

IndyCar (AI-Optimized + Swarmia)

Problem detection

After sprint/milestone, (missed sprint, spillover)

As patterns emerge (review time trending up)

Core dashboard vitals

Sprint burndown, velocity, backlog health, ticket status

PR cycle time stages, DORA, Investment Balance, Work Log, Focus Summary

Granularity

Sprint level ticket aggregates, Story points

Commits, Comments, PR stages (author → review → merge), developer level, agent-level

Quality / stability signals

Bug spikes after "Done", Post-sprint defect reports

DORA drift (lead time grows, deployment frequency drops, MTTR rises, CFR spikes)

Bottleneck detection

Growing "In Progress" column, long-lived tickets (visible after accumulation)

Real-time PR review latency, merge delays, stage-specific slowdowns

Engineering leader

Periodic check-ins, Manual report pulls

Pit-crew telemetry mindset, working agreements, proactive unblocking

Developer

Standup status updates, manual ticket updates to show progress

Automated nudges + notifications, focus on coding, actionable alerts



The Trifecta of Success



The AI Culture Shift: Strategic Decisions

Model 1: Interactive (Cursor)

Developer + AI working together in the IDE.

Curated System Prompts



Model 2: Semi-Attended (Claude Code)

Developer queues a task, handles planning/research, and supervises execution.

Guided Execution



Model 3: Fully Unattended (Agentic)

Autonomous agents pick up Jira ticket and complete the work end-to-end.

Jira Ticket:
"Add Payment Integration"



The Secret Sauce



From writers to orchestrators

Engineers haven't written code without AI in 12 months



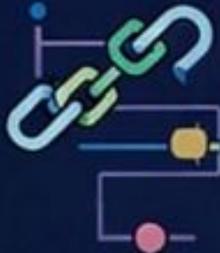
UAT Experts

Senior engineers are now taking on user acceptance testing and security audits



Engineer Autonomy

Autonomy is earned through business and technical domain mastery and enabled by org context



Prompting vs Hooks

Engineers don't "prompt harder". Instead the org creates context that drives agent behavior. Constrain AI with skills, hooks and guardrails so outputs stay consistent.

Org Shift Needed



The Engineer (Business + Tech Domain Mastery)



The Organization (Invests in providing the Context)



The Manager (Creates the Coach/Apprenticeship Environment)

It is even easier to be a good human with AI because mentorship and positive reinforcement deliver outsized ROI in an AI-augmented world.



AI impact

Last 30 days

Engineering AI assistant Any AI agent Any More filters...

vs Previous timeframe Group: AI used Average

Cycle time

Claude Code No AI Cycle time
3.3 days -35% 35 hours -21% 2.9 days -31%



Time to first review

Claude Code No AI Time to first review
20 hours -57% 15 hours +81% 18 hours -48%



Batch size

Claude Code No AI Batch size
749 -71% 896 -45% 784 -66%



Time in review

Claude Code No AI Review time
34 hours -49% 25 hours -37% 31 hours -47%



What it looks like in practice: last 30 days



AI experimentation is a journey

AI impact

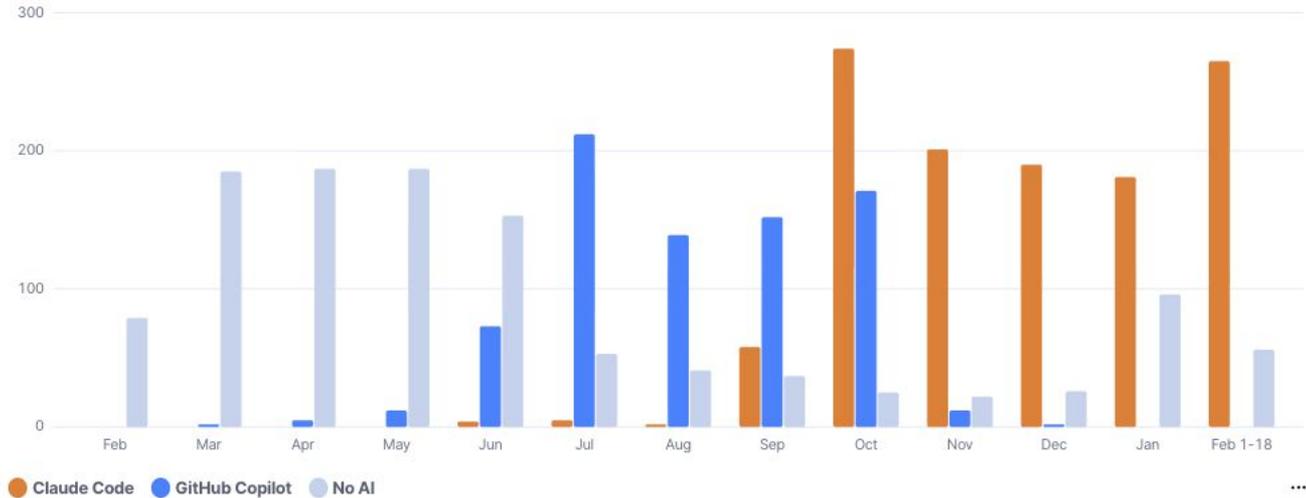
📅 Last 365 days ⓘ

📍 Engineering ▾ AI assistant Any ▾ AI agent Any ▾ More filters...

⇌ vs All selected PRs ▾ 📊 Group: AI used ▾ Σ Average ▾

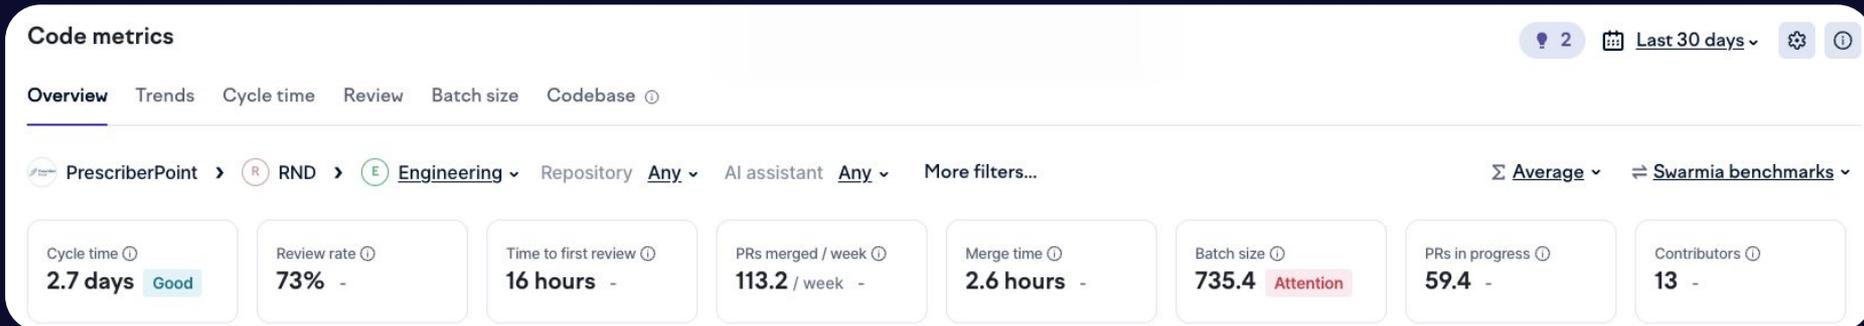
PRs merged ⓘ

Claude Code	1180
No AI	1147
GitHub Copilot	780
Total ⓘ	2873

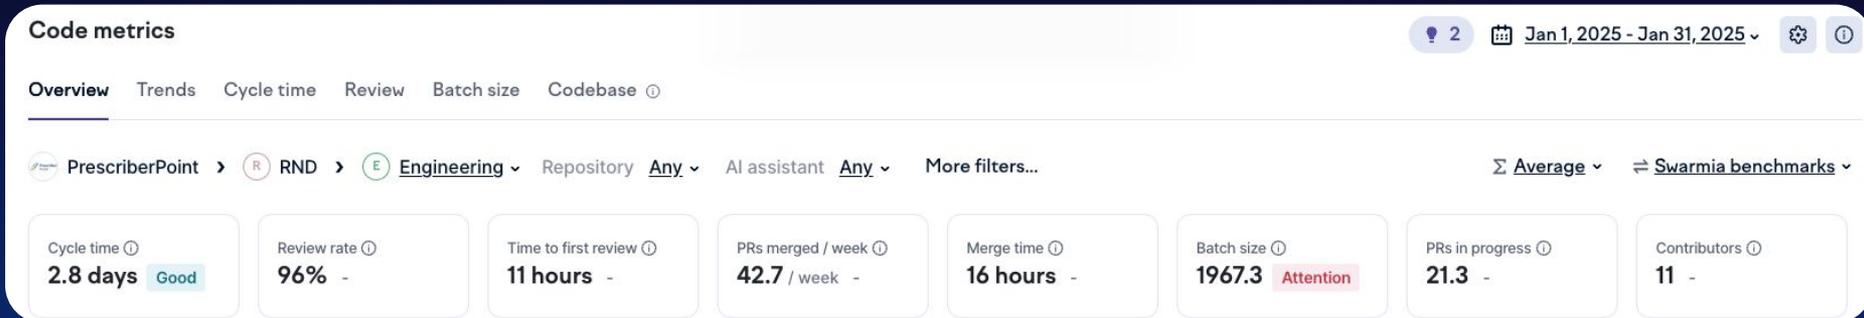


Now vs Baseline

Now



Baseline



Now vs Baseline

Now

Deployment frequency ⓘ

14 / week +91%

Change lead time ⓘ

14 days -32%

Time to deploy ⓘ

9.9 days -37%

Change failure rate (CFR) ⓘ

15% +5

Mean time to recovery (MTTR) ⓘ

5.7 days -55%

Baseline

Deployment frequency ⓘ

3.4 / week

Change lead time ⓘ

12 hours -5%

Time to deploy ⓘ

-

Change failure rate (CFR) ⓘ

13% -2

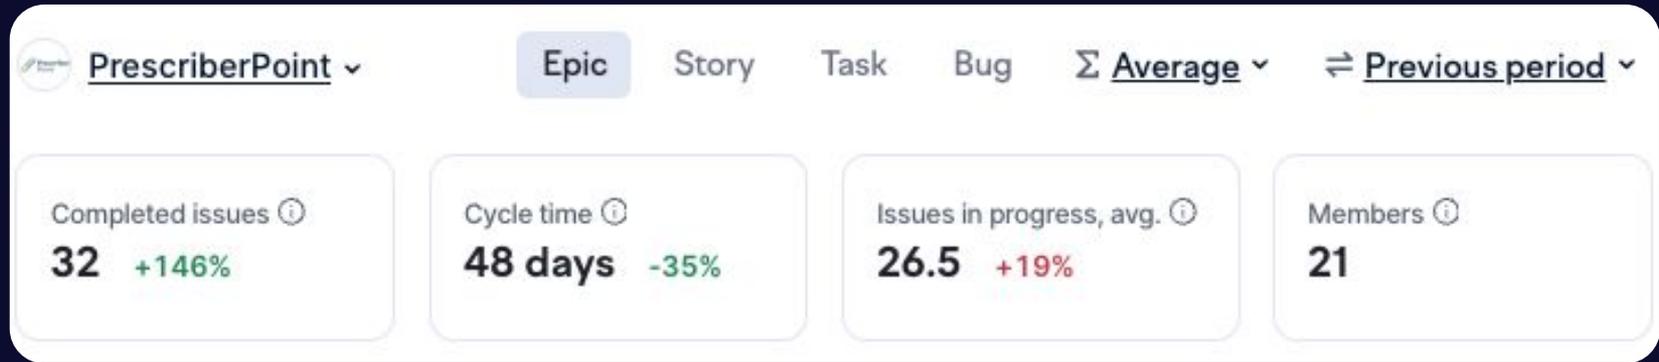
Mean time to recovery (MTTR) ⓘ

47 hours -14%

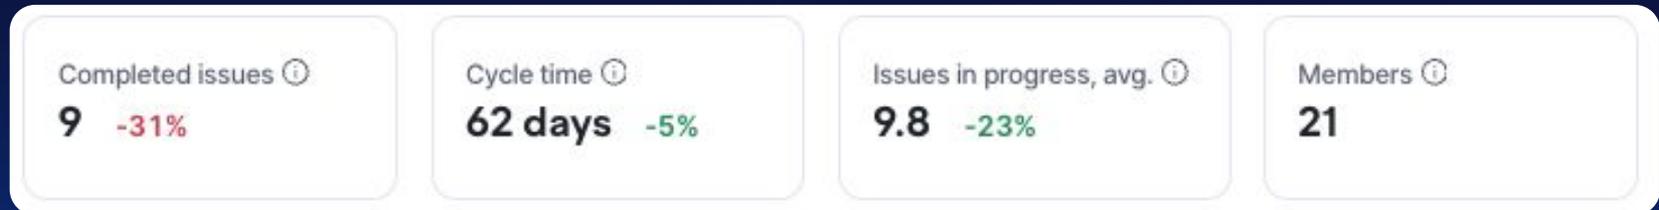


Now vs Baseline

Now



Baseline



Next Steps

- Fix your process and data hygiene first: You can't accelerate a broken foundation.
- Allow room to learn: Start with a small team to prove the method.
- Measure outcomes, not activity: Focus on features shipped and business value that was realized from those features, not lines of code.
- Starter AI workspace

QR code url: <https://github.com/theJohnHarbison/starter-ai-workspace>





 PrescriberPoint®

Recommended Reading

Learn more about how PrescriberPoint fixed their engineering fundamentals with the help of Swarmia before going fully agentic.



bit.ly/4bWpzh9



What tool are the metrics from?

Swarmia is a software engineering intelligence platform that helps you answer questions like: What does our engineering productivity look like with AI vs before? Is our AI investment paying off? Who are the most advanced AI users who could coach their peers?

Book a meeting
with us to learn
more



bit.ly/3MJZKXf

