

EXECUTIVE PRESENTATION • 2020

AI Initiative Wirk

Measured results from our own builds. A path to scale them across 160+ developers. The decisions that make it happen.

wirk

we create software

400h PROJECT

400h → 50h 8x faster

Effort reduction on our own pilot

DEV WORKFLOW

4h → 15min 16x faster

Maintenance screen build time

AT SCALE

160+ devs → multiplied

One standard, applied across the whole team



FASTER DELIVERY

How It Works for Developers

SCAFFOLD ARCHITECTURE

```
scaffold/
├── CLAUDE.md.tpl - project identity + key commands
├── Conventions - Things That Will Bite You
├── hooks/ - 14 Python scripts, one concern each
│   ├── guard_style.py - doc_checks.py - memory_extract.py
└── discovery/stack_markers.json - 8 stacks auto-detected
```

```
01 Developer opens any project
02
03 One-time /scaffold init deploys CLAUDE.md, hooks, and settings. Never repeated.
```

Session starts. CLAUDE.md auto-loads
Stack, conventions, gotchas, and key commands injected before the first prompt. No re-explaining the codebase.

Claude generates convention-aware code
kebab-case files, PascalCase classes, correct API shape, right branch naming. All from context, not instruction.

400h Project **Efficiency**

THE SCENARIO

A development team quotes a large project: 400 hours estimated. With the AI scaffold applied end-to-end, initial UAT delivery came in at 50 hours.

Same developer, same quality standard. The scaffold handled spec generation, code scaffolding, test generation, and documentation at each stage. At Wirk's scale, this multiplier across 160+ developers translates to thousands of additional delivery hours per quarter without additional headcount.

WHAT THIS MEANS FOR WIRK

Every large project in the pipeline carries the same opportunity. More capacity per developer means faster client delivery, more projects in parallel, and margin improvement without scaling headcount.

BEFORE

400h

MANUAL WORKFLOW

AFTER (SCAFFOLD APPLIED)

50h

AI SCAFFOLD APPLIED

THE RESULT

87.5%

EFFORT REDUCED

Same developer. Same quality. 8x faster. UAT-ready output, not a prototype.

Legacy Modernisation **Acceleration**

THE SCENARIO

A client has a legacy codebase: undocumented, tightly coupled, built over years. Understanding it, mapping dependencies, and producing a migration plan takes weeks manually.

AI-assisted legacy analysis reads the codebase, maps dependencies, identifies high-risk areas, and generates a structured migration plan with effort estimates. The analysis that took two weeks takes two days. The plan is documented, reviewable, and client-shareable from the first session.

WIRK'S SERVICE DIFFERENTIATOR

Legacy modernisation is already a Wirk service. AI-accelerated analysis makes this faster, more accurate, and more defensible to the client. It also produces the documentation that protects Wirk's team during the migration itself.

01 / / /

AI codebase analysis

Reads legacy codebase, maps dependencies, identifies coupling and risk areas. Two weeks of manual analysis in two days.

02 / / /

Migration plan generation

Structured migration plan with effort estimates, risk ratings, and recommended sequence. Client-shareable from day one.

03 / / /

Documentation as byproduct

Every migration step documented automatically. Protects the team and the client throughout execution. No knowledge locked in one person.

04 / / /

Faster scoping, faster delivery

More accurate scoping means more accurate quotes. Faster analysis means faster project starts. Both directly improve margin and client confidence.

PM Workflow & ClickUp Pipeline

THE SCENARIO

A PM writes one prompt. Three minutes later: a working, on-brand, interactive mockup exists. The use case is built. The technical spec is written. A ClickUp task is created and tracked.

Without the scaffold, this PM was spending 3+ days per deliverable with inconsistent results. With it, the entire brief-to-tracked-work-item pipeline runs in one session. The developer starts from a verified, documented task. No manual handoff. No lost context.

THE SCAFFOLD IS THE DIFFERENCE

The mockup generator reads the design system, follows composition recipes, applies tokens and components, then verifies against a live replica. The output isn't "an idea." It's a production-ready screen that passes verification. That's what makes the ClickUp task trustworthy enough to send to development immediately.

01 / / /

PM writes one prompt

Describes the screen or feature. No wireframe, no spec document, no design brief needed upfront.

02 / / /

Scaffold generates verified mockup

Reads design-system.yaml, follows recipe rules, applies correct tokens and components, verifies against live replica. Output is production-ready.

03 / / /

Use Case Builder produces spec

Structured use case + full technical spec for developers. Same format every time. No variation, no gaps.

04 / / /

ClickUp task created automatically

Work item created with spec attached and context preserved. Development starts from a tracked, verified work item. No manual handoff required.

Support Operations **at Scale**

THE SCENARIO

Wirk manages 180+ applications and 1,000+ devices across manufacturing and dealer networks, 24/7. Incident triage, root cause analysis, and documentation are constant workloads.

AI-assisted support operations triage incoming incidents, suggest probable root causes based on historical patterns, generate structured incident reports automatically, and flag recurring issues that indicate systemic problems. Engineers spend time resolving issues, not administering them.

THE ITIL ADVANTAGE

Wirk's ITIL-certified three-tier support model is already structured for AI augmentation. AI triage sits at the first tier, routing incidents with context before a human engineer touches them. Resolution time drops. Pattern detection improves. The same team covers more ground.

01 / / /

AI incident triage

Incoming incidents classified, prioritised, and routed with suggested probable causes. Engineers start with context, not a blank ticket.

02 / / /

Pattern detection

AI surfaces recurring incidents that indicate systemic issues. Problems flagged before they escalate to major incidents or client-visible failures.

03 / / /

Automated documentation

Incident reports generated automatically from resolution notes. Compliance documentation maintained without manual effort. Audit trail always current.

04 / / /

Same team, more coverage

AI handles first-pass triage and documentation. Engineers apply judgment where it's needed. No headcount increase required to expand coverage.

Our AI Initiative: What We're Building

01 ///

Education

Structured AI literacy for every developer: tools, prompting, and responsible use. Built on the failure modes already documented so teams skip the learning curve.

02 ///

Standardisation

Defined AI workflows and proven methods. Codified processes that replicate the 15-minute builds and 50-hour projects across every team, every client, every engagement.

03 ///

Policy & Governance

Clear policies on what tools to use, how data is handled, and what outputs require validation. Especially important for manufacturing and industrial clients with strict data requirements.

04 ///

Measurement

Tracking time saved, quality improvements, and ROI across teams and client engagements. Makes the business case visible and reportable at every leadership review.

05 ///

Team & Resourcing

One person owns standardisation, iteration, and feedback. That ownership model, applied across Wirk's teams and divisions, is how this scales beyond individual effort.

06 ///


Executive Alignment

Three asks: public leadership commitment, approval for the short-term productivity buffer during upskilling, and policy sign-off. These are what move this from individual effort to company capability.

Understanding the Adoption Curve

There is a learning dip. We are planning for it.

When developers begin adopting AI, initial productivity drops 20–30% as they map the learning curve. This is expected and temporary. The real budget cost is billability, not cash. The scaffold shortens that dip significantly, because the failure modes are already documented. Teams that follow it recover faster and stabilise higher.



WHY THE FIELD IS ACCELERATING

LLM inference costs fell 40x per year (Epoch AI). The same AI capability costs 280x less than two years ago (Stanford AI Index 2025). Every month of delay means competitors access better tools at lower cost.

01 / / /

The Dip Is Temporary

Productivity drops 20–30% as developers adjust. Those who persist (especially with scaffolded guidance) recover and exceed their original output.

02 / / /

The Scaffold Shortens It

The failure modes are already mapped. The scaffold gives developers a structured path through the learning curve, not trial and error.

03 / / /

Stagger the Rollout

One team at a time. At Wirk's scale, this keeps client delivery stable while the standard is refined and proven before broader rollout.

What We're Asking For

01

Budget Approval

The real budget is billable time, not direct cash. People's billability dips temporarily as they upskill. Approval means accepting a short-term billability buffer while teams clear the adoption curve. The payback begins immediately and compounds with every project.

02

Executive Sponsorship

Visible leadership endorsement to signal this is a company priority, not a team experiment. At 160+ developers, a top-down signal is what creates the conditions for consistent adoption.

03

Dedicated Headcount

One person with clear ownership to drive standards, training, and continuous improvement. This moves the initiative from person-dependent to institution-dependent across all of Wirk's teams and client engagements.

04

Policy & Governance Sign-off

Executive approval of an AI usage policy, formalising what is encouraged, permitted, and required. Essential for Wirk's manufacturing and industrial clients who operate under strict data and compliance requirements.

Addressing Your Concerns

"Our clients need consistency and precision. AI introduces risk."

The scaffold is precisely about removing that risk. It doesn't replace discipline. It encodes it. Every output passes a live-replica verification before it's considered done. Manufacturing clients get the same quality standard they always have, delivered faster and with documented context.

"If this is working, why isn't it spreading already?"

Results don't spread on their own. They spread when someone structures the evidence clearly enough that a decision becomes easy. You're experiencing that right now. The initiative formalises that process so it isn't dependent on one person doing it. That's how individual results become company results.

THE DATA IS CLEAR · THE RESULTS ARE REAL · THE RISK OF DELAY IS MEASURABLE

We move from pilots to practice together.

Four decisions. One direction. Secure our competitive position.

- Approve the budget to scale our AI initiative
- Provide executive sponsorship to signal company-wide commitment
- Authorise dedicated headcount to own and drive AI standards
- Sign off on the AI governance policy to protect the business and our clients

WIRK · AI INITIATIVE · 2026

