

fortégroup Xceptor

Beyond AI Pilots: Operationalising AI in Software Delivery

March 10, 2026

fortégroup



Alex Lukashevich
Chief AI Officer
Forte Group

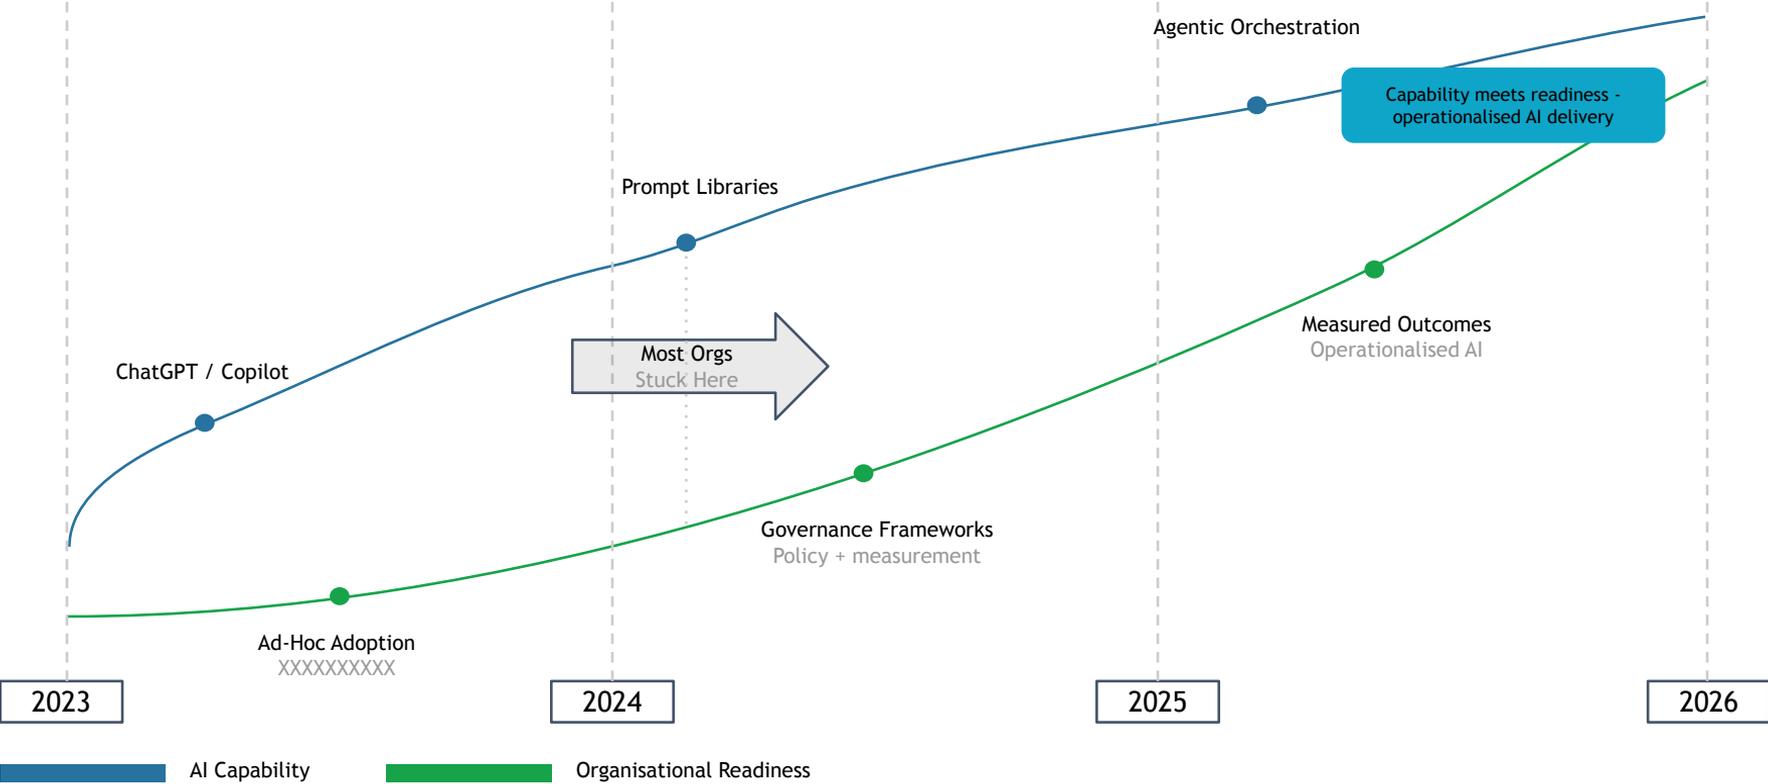
Xceptor



Mike Kinloch
SVP Engineering
Xceptor

The Convergence Moment

Why AI in Delivery is Finally Working



Xceptor - Our mission

We **empower** business users within **financial institutions** to build **automated processes** that deliver **trusted data**.

Xceptor



The engine powering financial markets

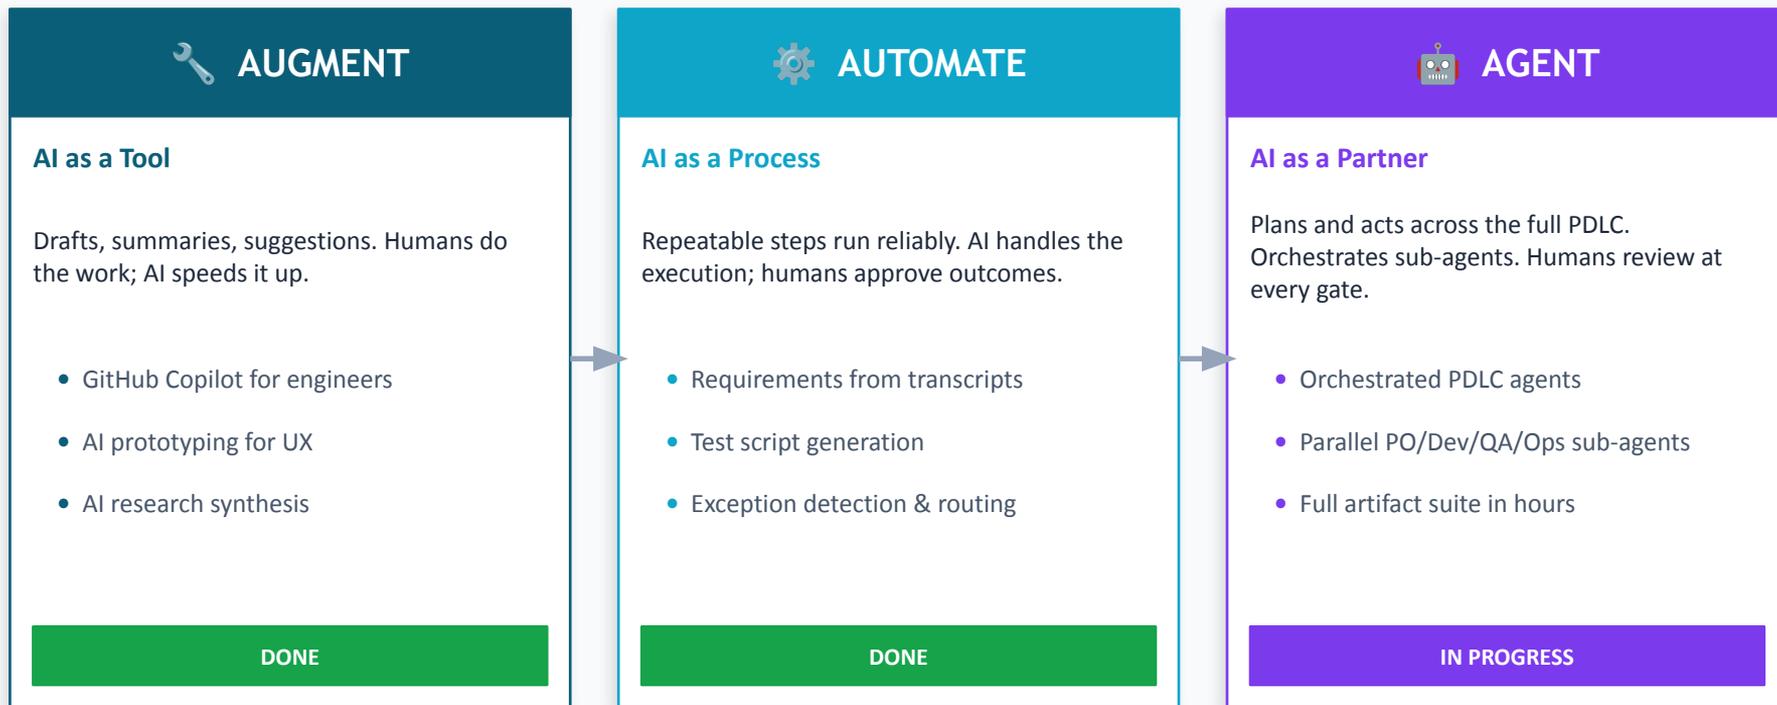
Streamline pre- and post-trade workflows via data automation & process digitization

PURPOSE-BUILT PRODUCTS



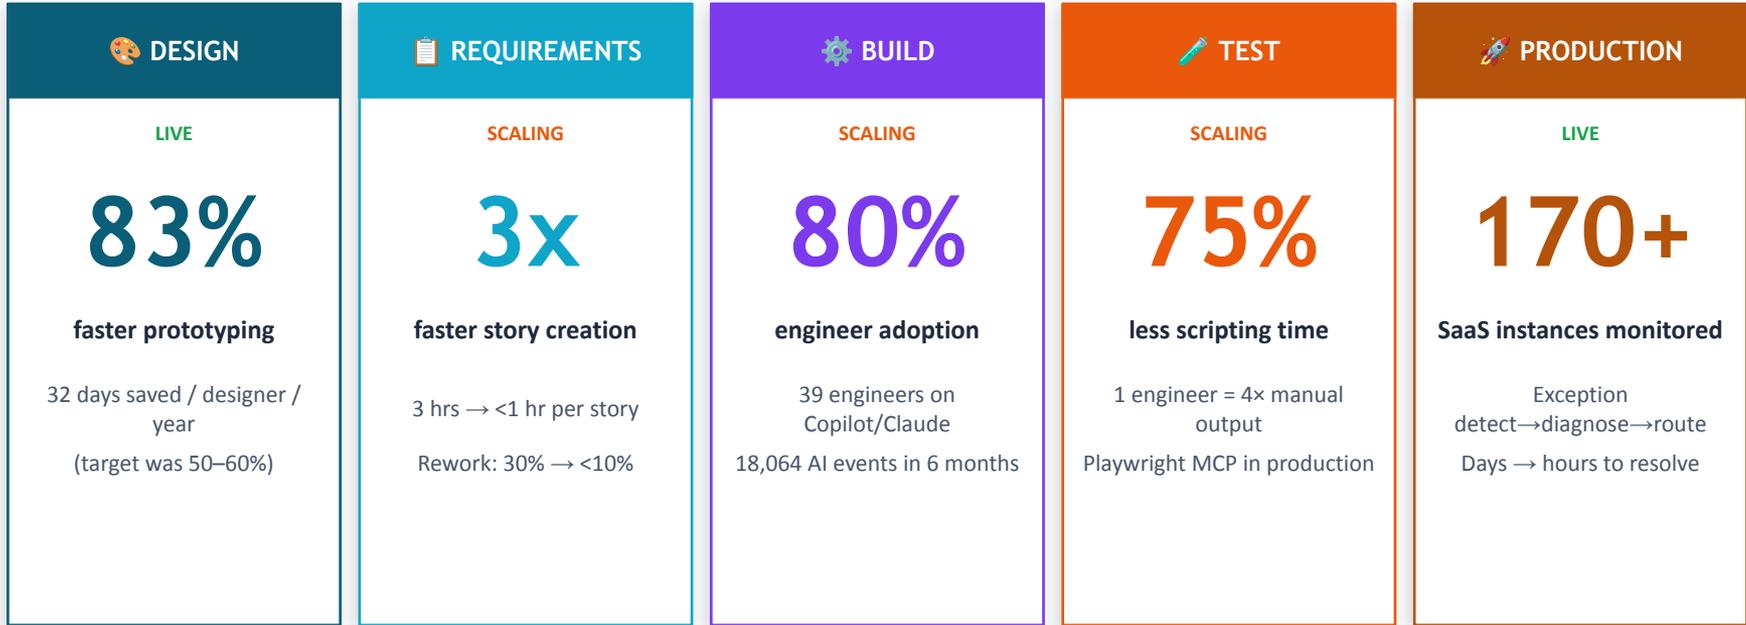
From AI-Assisted to Agent-Driven

A deliberate three-stage journey - start small, measure impact, then scale



AI Embedded Across All Five PDLC Phases

Five active areas delivering measurable results



51.7% of all AI usage across Xceptor is in Engineering & Product - we're leading from the front

Real Challenges – and How We Overcame Them

The honest picture: what went wrong, what we learned, and how we adapted



AI model quality varied wildly

GPT-4 generated brittle test scripts; Copilot struggled with our custom Selenium setup.



HOW WE SOLVED IT

Switched to Claude Opus for test generation. Moved from Selenium to Playwright MCP - AI-native and well-documented.

OUTCOME

75% scripting time reduction achieved



Building first AI agent

Our SaaS monitoring agent had to make real decisions across 170+ instances - false positive rates and reliability were a genuine challenge.



HOW WE SOLVED IT

Read-only mode first. Confidence thresholds before routing. Human sign-off on novel exceptions. Incremental rollout.

OUTCOME

Exception triage: days → hours; zero missed P1 incidents



Adoption wasn't automatic

Engineers resisted AI tooling when it didn't fit their workflow. Mandating Copilot without enablement failed silently.



HOW WE SOLVED IT

Champions programme, role-specific playbooks, and measuring outcomes (not just install counts). 80% adoption came from pull, not push.

OUTCOME

39/49 engineers active; 18k+ AI events in 6 months

The AI-Native PDLC

A central agent ecosystem where AI builds and humans review - governed, auditable, and directed by our people

01

Lightning Speed

AI is the primary builder. Humans review, guide, and steer. We move at a fundamentally different pace — features in hours, not weeks.

02

Robust Control & Governance

Every agent enforces quality by design. Governable, tunable, and auditable at every decision gate. No black boxes — ever.

03

Central Agent Ecosystem

One unified platform. POs, Devs, QA, and Ops all interact with purpose-built agents. Best-in-class UX for every role.

"AI builds at speed. Humans stay in control."

Orchestrator & Sub-Agents: Parallel Execution

A single orchestrator agent receives a feature request and instantly coordinates specialised sub-agents running in parallel.

-  Orchestrator plans, delegates, and collects all outputs
-  PO, Dev, QA, DevOps agents run simultaneously
-  Human review gate before any output is promoted
-  Full artifact suite: stories, code, tests, deploy plan, release notes
-  Every decision auditable and human can override at any gate

AI PDLC Agent Orchestration

How sub-agents execute your PDLC in parallel

with human approval at every gate

ROLE EVOLUTION

How PDLC Roles Transform

The AI PDLC changes what our people spend their time on - from doing to directing.

TODAY	TOMORROW
-------	----------

Product Owner	Writing features and user stories manually (3h+)	Reviewing & steering agent feature and story output
----------------------	--	--

Developer	Writing code & docs	Reviewing PRs, architecture decisions
------------------	---------------------	--

QA Engineer	Manual test scripting (4x slower)	Approving AI test suites, edge cases
--------------------	-----------------------------------	---

Ops / DevOps	Triaging 100s of exceptions daily	Reviewing AI insights, approving deploys
---------------------	-----------------------------------	---

One person can now span what previously required three - teams shrink, scope expands

The 6-Week MVP: Proving It in Production

Goal: full connector from requirement to production-ready in under 1 day — every artifact, every gate, every review

SUCCESS CRITERIA

- ✓ Production-ready connector delivered in <1 day
- ✓ All artifacts present: stories, arch docs, code, tests, release notes
- ✓ Every output reviewed and approved by a human
- ✓ Reusable agent framework any team can pick up
- ✓ Training playbooks so adoption doesn't need hand-holding

WHAT WE'LL MEASURE

-  Time per PDLC phase vs. baseline
-  Severity bugs found in review / rings
-  Team satisfaction (DX pulse)

AGENT ECOSYSTEM CAPABILITIES

PO Agent: Requirements, user stories, acceptance criteria

Dev Agent: Architecture docs, code, PR, technical design

QA Agent: Test strategy, test cases, test data

DevOps Agent: Security review, release summaries, rollback plans

Cross-cutting: Assisted code review, issue triage, onboarding playbooks

Governance: Policy guardrails, audit trail, human override at every gate

Week 1–2

Discovery & setup

Week 3–4

Build & integrate

Week 5

Validate & test

Week 6

Deploy & measure

What the AI-Native PDLC Means for Xceptor

 Move Faster

100×

delivery velocity potential

Features that take weeks of sequential handoffs will be delivered in hours. Parallel agent execution reduces wait time at every boundary.

 Higher Quality

< 10%

rework rate (down from 30%)

AI-generated artifacts are consistent, source-linked, and reviewed at every gate. Fewer defects escape to production - and we catch them faster when they do.

 Stay Competitive

3×

more scope per engineer

We grow capability without growing headcount linearly. Our people operate at higher leverage - focused on strategy, creativity, and judgement. Not formatting, scripting, or triaging.

We are building this now!

fortégrouP

Xceptor

Thank you!