

AI Production Readiness Sprint

Your AI Works. Approval is the part that's stuck

A verdict & One live control — not a report, not the full-build

Three weeks, one AI system. We find what's blocking production approval and switch on one real control



TRUSTED BY ORGANIZATIONS ACROSS

Financial Services · Insurance · Healthcare · Technology & SaaS · Public Sector

SOC 2 Type II · ISO 27001 · GDPR · Sector Frameworks (PCI-DSS, HIPAA & more) on request · ~14 delivered AI engagements



We take AI the last, hardest mile — on an outcome basis

We take AI from pilot to production, in every shape, with senior engineers who are AI-augmented inside a well-governed process. Backed by our own in-house AI lab, and partnerships with NVIDIA, Microsoft and AWS.



Swiss-owned, built by the majors

A global AI engineering group, Swiss-owned and delivering from Vietnam, with offices in the UK, Singapore and Riyadh. Built by people from Luxoft, PwC and other global technology and consulting firms.



Highly regulated industries

100+ clients and projects delivered across EU, APAC, Middle East and US markets — strong experience in highly regulated industries.

Why Now?

AI pilots are no longer the challenge — enterprise adoption is. More customers, regulators and internal governance teams now expect **evidence** that **AI systems are secure, governed and production-ready** before they approve deployment. Companies that solve this gap earlier ship AI features faster and win enterprise deals sooner.

THE CHALLENGES

Pilots that struggle to reach production

- ❗ The pilot has been demoed twice, and the third request came from someone who should be approving it, not watching it again.
- ❗ Security or risk asked for evidence that doesn't exist yet.
- ❗ Leadership wants a yes/no answer this quarter, not a new workstream.
- ❗ The team that built it can't grade it — they need an outside verdict, not another internal opinion.
- ❗ A go-live date is on the roadmap with no one accountable for whether it will pass review.
- ❗ Nobody in the room can say what would need to be true for this to get approved.

OUR SOLUTION

What if you had a straight answer in three weeks — not a maybe?

The Sprint is a fixed-scope, fixed-price engagement on one AI application. We map it, measure it against a recognised framework, and switch on one control that actually runs in your environment — not a recommendation to build one later.

3 weeks

kick-off to a decision your board can act on

1 live control

running in your system before we leave

NIST AI RMF

scored in the language your auditors accept

Fixed price

approvable by one owner, no transformation budget

Software tells you what's happening. Our engineers make the system production-ready — and prove it.



We build, not just advise

The people who assess the system are the people who fix it.



AI-amplified delivery

We rebuilt our lifecycle around human-AI collaboration — we think, architect and test at a different pace.



Journey together

Diagnose, then build, then operate. Extends directly into Pilot-to-Production and the Control Desk.

How to Start

Three weeks, three moves — inside your perimeter

Look

WEEK 1

Map the system, its data and its access.
Agree the questions that matter.

YOU RECEIVE

System & dataflow map / dependency list.

Measure

WEEK 2

Run reliability, fairness and explainability tests. Score readiness on the NIST AI RMF.
Choose the control.

YOU RECEIVE

Evaluation baseline / readiness score / control spec.

Fix & Decide

WEEK 3

Switch on the control. Finalise the gap register, the 90-day plan and the go/no-go memo.

YOU RECEIVE

Live control / gap register / 90-day plan / go-no-go memo.

The senior POD on this engagement



Alexander Nevolin
Head of Delivery Manager
25+ years of experience in software engineering and technology leadership



Yurii Lozinskyi
Chief Product Officer
25+ years driving Digital Transformation across Healthcare, Insurance, and Finance



Dr. Toan Nguyen Huy
Chief Technology Officer
PhD, Computer Engineering. 8+ years in applied AI, CV and LLMs



M.Sc. Thang Nguyen
AI/Computer Vision Tech Lead
AI specialist with experience in end-to-end vision systems



M.Sc. Tuyen Dinh
Senior AI Engineer
MSc, KAIST. 5+ years in ML, deep learning, CV and NLP



Quan Tran
Senior AI Engineer
6+ years in enterprise AI, banking and financial automation



Viet Anh Dang
Senior AI Engineer
5+ years in enterprise AI, telecom and insurance

SUCCESS STORY · DOCUMENT VALIDATION AT SCALE

An LLM system that validates high-volume documents



93%

Validation Accuracy

PROBLEM

Documents were validated by hand. Slow, inconsistent, and errors reached downstream systems.

-70%

Human Errors

WHAT WE BUILT

An LLM system that reads and validates documents, with the evaluation baseline and one live control needed to run it in production.

-40%

Manual Review

IN PRODUCTION

Measured against a held-out set, with humans kept on the cases that need judgement.

First proven in insurance KYC; the same pattern — high-volume document or data review with humans kept on the judgement calls — repeats across lending, healthcare intake and claims. Figures are client-reported.

Pricing Model

Fixed scope, fixed price, about three weeks — approvable by one owner, without a transformation budget. We price it against your system on the scoping call.

Diagnose → Build → Operate

This Sprint tells you if it's safe to ship. Pilot-to-Production takes it live. The Control Desk keeps it healthy.

NEXT STEP

Book a 30-minute AI Production Readiness conversation

Bring one AI system that's stuck. We'll identify what it is, where it's blocked, who needs to approve it and what evidence is missing. If it isn't a fit, we'll say so on the call.

Vladimir Vereschagin
CEO
vladimir@smartdev.com

Pavel Veselovskiy
Senior Business Development Director, AI Solutions
pavel.veselovskiy@smartdev.com

Minh Duong Vu
Business Development Manager
minh.duongvu@smartdev.com