

The pilot worked. The rollout didn't?

An 8–10 week engagement that turns a validated concept into a production-ready MVP — built to absorb into how your business actually runs.

8–10 wks

Concept to working MVP
Iterative, validated

3

Phases
Design → configure → implement

In prod

An operational use case
Deployed, not a demo

Day 1

Responsible AI built in
Governance, not retrofit

The Problem

Pilots everywhere. Production nowhere.

Demos impress in the lab and die on the way to production — no real environment, no data pipeline, no path to scale. The MVP is where AI stops being a slide and starts absorbing into your operating model.

What It Does

- **Build in your environment**
Configure infrastructure, MLOps, and AI frameworks on your own stack.
- **Develop iteratively**
Two-week cycles, continuous validation, pre-trained models where they fit.
- **Wire in your data**
Automated ingestion from real sources, transformed into usable models.
- **Ship to production**
Deployed MVP, technical blueprint, runbooks, and knowledge transfer.

POC / MVP · how the 8–10 weeks work

ENGAGEMENT PHASES

Phase 1

Design & planning
2 weeks

- Resource requirement planning
- Proposed solution architecture review
- Determine optimal hardware & software configurations.

Phase 2

Configuration & setup
2 weeks

- Install and configure AI POD for MVP scope
- Onboard and integrate with existing IT
- Configure dev environment with MLOps and AI frameworks
- Deploy and configure data-pipeline tools

Phase 3

MVP implementation
6 weeks

- Two-week iterative development cycles.
- Model development and training; pre-trained models where they fit.
- Automate data ingestion; transform into analytical models.
- Develop end-user experience; optimize for scale.
- Deploy to production; knowledge transfer and closure.

DELIVERABLES

Build artifacts

- Working prototype.
- AI POD configuration.
- MLOps configuration.

Documentation

- Technical architecture blueprint.
- Implementation guides.
- Maintenance & support procedures.

WHY ONSTAK

Engineers, not slideware

- 75% AI architects and engineers — the experts of the experts.

Built and run as one

- We integrate. We are not an integrator. One team, no handoff gap.

Outcome-based. Human first.

- We start with you, then deliver aligned to your needs.

Duration: **8–10 weeks** · Output: **an operational AI use case in production, ready to scale.**