imobisoft

The Lean AI roadmap

Smarter business, smaller budget

A practical guide for UK mid-market leaders working with legacy systems

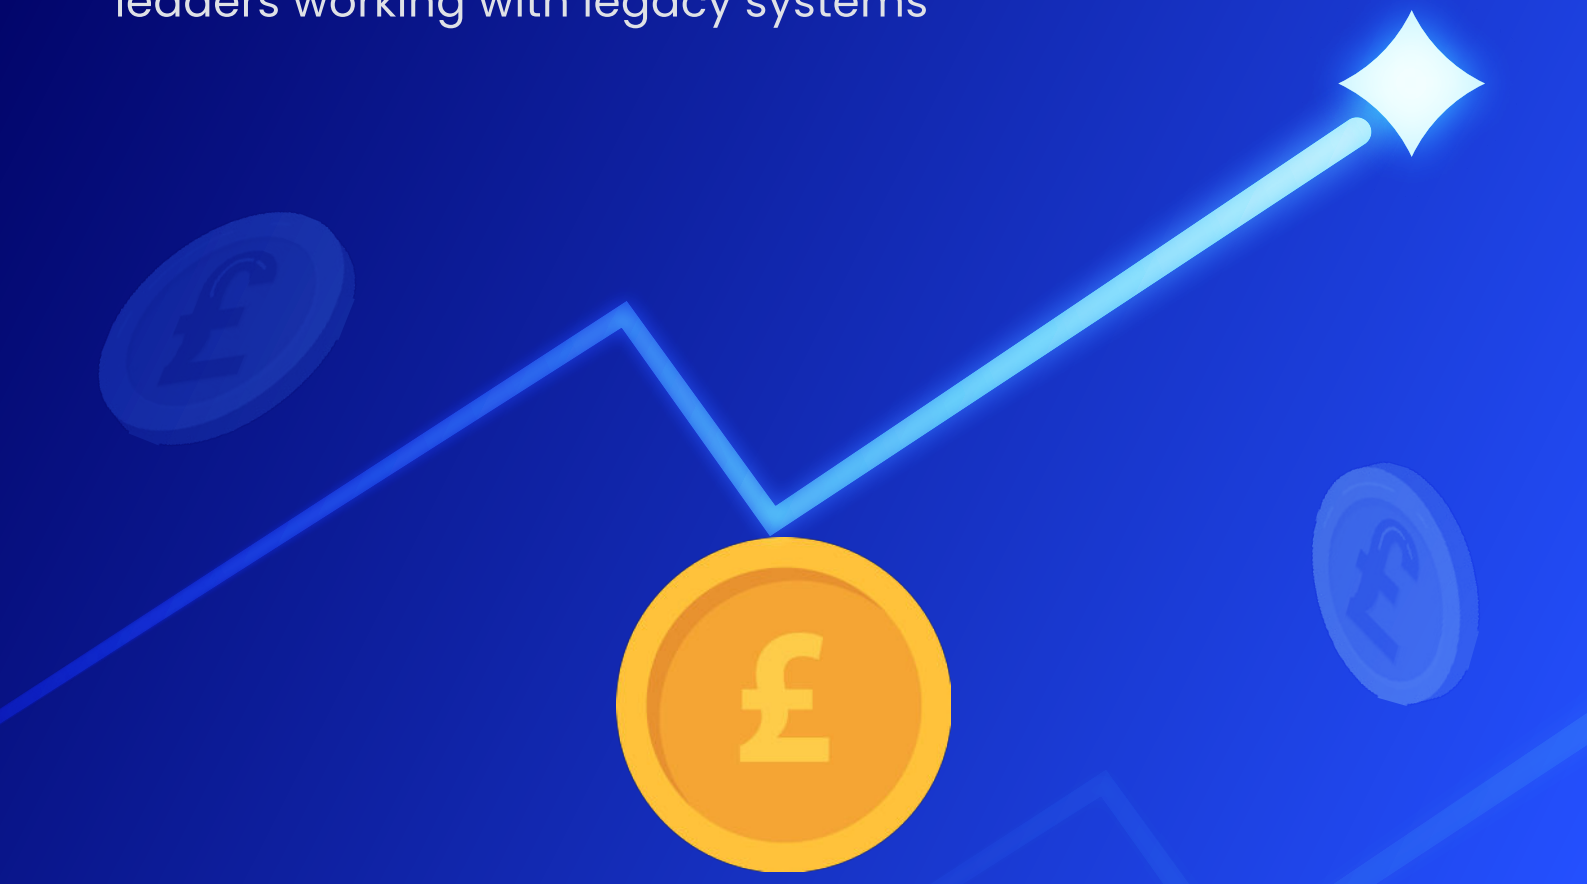

By Imobisoft Ltd

The UK mid-market is under pressure. You're expected to innovate to stay competitive, but you don't have the budget, time, or risk appetite of a Silicon Valley giant.

Yet the dominant AI narrative says the same thing: transform everything. Replace legacy systems. Rebuild your data stack. Hire expensive specialists.

We disagree!

> *"The biggest mistake companies make is thinking they have to rebuild everything from scratch."*
>
> *This is why 42% of companies abandon their AI initiatives.*
>
> *(Source: S&P Global Market Intelligence)*

And the problem is growing. Gartner predicts that by 2027, over 40% of Agentic AI projects will be cancelled, largely due to escalating costs and unclear business value.

In our experience, these failures follow a familiar pattern. Costs spike when organisations attempt a single, "big bang" AI transformation overhauling complex data architectures instead of working with the systems and workflows they already have.

This is exactly what Lean AI is designed to avoid.

Lean AI tackles AI failure at its root by surgically targeting data waste, not rebuilding everything. The result is faster deployment, clearer outcomes, and up to a 30% reduction in operational costs without disrupting the business.

That's why Imobisoft advocates a lean, disciplined approach to AI adoption. Research and real-world delivery both show that companies succeed when they stop treating AI as a science experiment and start treating it as a business improvement programme.

We believe AI shouldn't be a massive, risky overhaul. It should be a practical tool that fixes specific problems, saves money, and pays for itself quickly.

The strategy

The problem

Rigidity in old systems: Your business relies on systems such as invoicing, payroll, and/or logistics. They are reliable, but they are old. Replacing them is risky, expensive, and takes years.

The solution

The "smart wrapper": You don't need to delete your old software to get new results. As a pragmatic AI development company, we use a technique known as the Strangler Fig pattern (or "The Wrapper"). This allows us to inject intelligence into your business without rewriting a single line of your legacy code.

Think of it like adding a smart thermostat to an older house. You don't knock down the walls or replace the wiring. Instead, you add a modern control layer that connects to what's already there, optimises how everything runs, and can be upgraded over time without disrupting daily life.

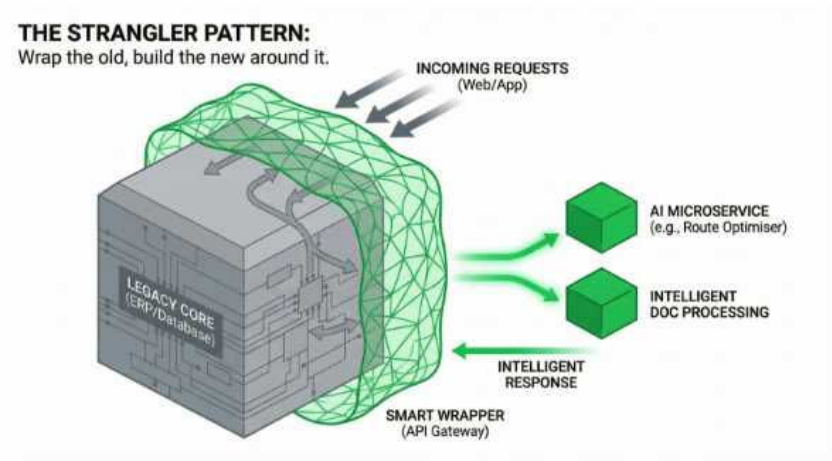

The Wrapper protects your core data while enabling modern features.

What is the Strangler Fig Technique good for:

- **Keep your core safe:** Your main business data stays untouched.
- **Add intelligence:** We build a "smart layer" that talks to your old system.
- **Move fast:** You can validate solutions in weeks, not years.

What the Wrapper isn't

The Wrapper is powerful, but on its own, it solves nothing. Without a clear target, it becomes just another delivery mechanism.

Wrapping the wrong process simply automates inefficiency faster. Wrapping the right one removes friction and pays for itself.

So the real question isn't how you add AI. It's where you add it.

Choosing the right target: eliminating waste before adding intelligence

Before we wrap anything, we apply a simple discipline: identify the work that consumes time and money without creating value.

We don't start with models, platforms, or architecture diagrams. We start by looking for friction, manual handoffs, duplicated effort, idle capacity, and brittle processes that slow the business down.

To do this consistently, we use principles from **Lean Six Sigma.** When we apply this lens to AI adoption, the same patterns appear again and again, regardless of industry or technology stack.

These patterns fall into three forms of waste that derail AI projects.

The three wastes that derail AI adoption

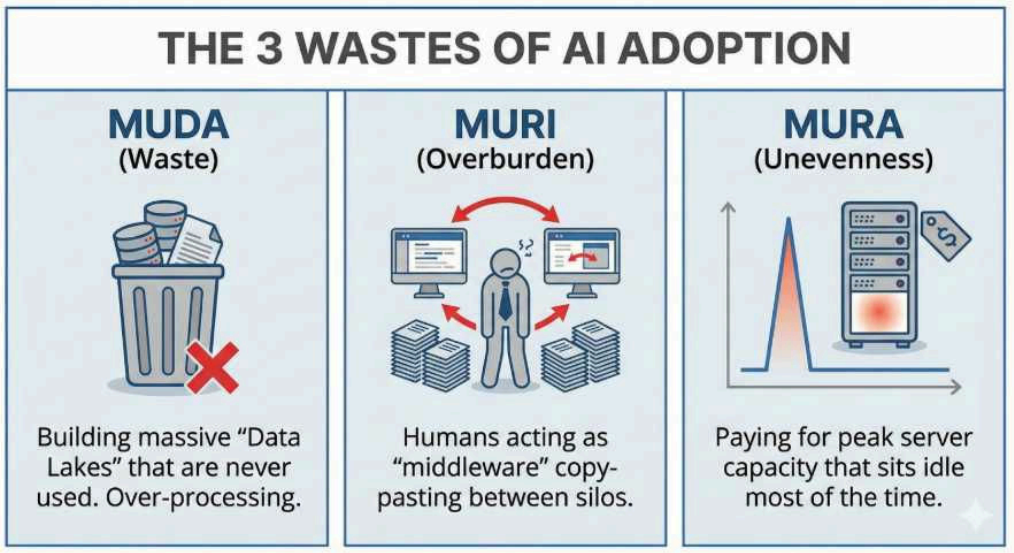

1. **The infrastructure trap**

   (Lean term: Muda — over-processing)

   The trap:

   The mistaken belief that you must achieve "Data Nirvana" i.e., cleaning all historical data before deploying a single algorithm.

   This turns AI into a massive capital project before it has proven any value.

   The lean fix:

   Just-in-time data.

   Value is only created when data is used. We clean only the data required for the immediate pilot. What looks like a multi-year transformation becomes a focused, low-risk experiment.

2 The “human middleware” problem

(Lean term: Muri — overburden)

The trap:

Legacy systems don’t communicate, so people fill the gaps. Staff manually transfer data between emails, PDFs, spreadsheets, and internal systems.

These handoffs become bottlenecks, slow down operations, and introduce errors.

The lean fix:

Intelligent Wrappers.

This is where the Wrapper earns its keep. We build a secure layer that reads documents, extracts data, and injects it directly into your existing systems without replacing them.

Your team stops acting as middleware.

The Safety Valve

"How do I know the AI won't make mistakes?"

We use **confidence scores** with **Human-in-the-Loop** review to catch uncertainty before mistakes happen.

High Confidence **(>95%)**

The AI processes the task automatically.

Low Confidence **(<95%)**

The AI flags the specific item for your team to review.

Result: You get automation speed without sacrificing control.

3 The capacity trap

(Lean term: Mura — unevenness)

The trap:

Traditional infrastructure forces you to pay for peak capacity all the time, like leaving the lights on in an empty office just in case someone walks in.

You fund idle systems to handle occasional spikes.

The lean fix:

On-demand processing.

Using serverless architecture, systems activate only when work arrives, process it, and shut down immediately. Costs align directly with usage and value.

But does this philosophy survive contact with reality?

Implementing our learnings

We tested these principles in two of the UK's most rigid, regulated sectors: Health care and Energy. The results were not just incremental improvements, they were transformative.

1 NHS (University Hospital North Midlands)

The friction Managing "Idiopathic Pulmonary Fibrosis" (IPF) patients required doctors to rely on infrequent, paper-based updates, meaning they often missed critical symptoms between hospital visits.

The lean fix We built "My IPF," a digital platform that connects patients directly to clinicians via a simple web app. Patients log symptoms daily, and automation flags risky trends immediately.

The result: A decade-old manual clinical pathway was fully digitised in just 6 months, allowing for proactive care that keeps patients safer.

2 E (Gas and Electricity) Ltd

The friction To stay competitive, E's team had to manually analyse massive competitor tariff files released by regulators every two weeks. The data was inconsistent and required days of manual cleaning and spreadsheet work every cycle.

The lean fix Imobisoft built a Machine Learning model to instantly ingest, clean, and analyse the competitor data.

The result: The process that used to take days now takes just seconds with data accuracy. This allowed E to react to market price changes instantly, without hiring extra admin staff.

Achieving these results required a calculated strategy, one grounded in a simple philosophy.

" What made the difference wasn't smarter models, but removing the unnecessary steps around them. Once we stopped moving data by hand, cleaned only what the decision needed, and used confidence thresholds to control automation, the systems became reliable enough to run in live, regulated settings."

Atif Syed, Chief Technology Officer, Imobisoft

But removing operational friction must never create legal friction. To ensure your AI is a compliant asset rather than a liability, you need rigorous safeguards.

We believe every AI partner should be able to prove they are safe. Use this framework to audit any vendor you consider starting with.

The "Safe Partner" framework

Use this litmus test to expose the gaps in a vendor's compliance strategy. If they stumble on any of these five points, they are not ready for a regulated environment.

Data sovereignty check:

Ask: "Can you guarantee my data never leaves the UK/EU?"

Why: Post-Brexit, data residency is a legal minefield. If they use US-based servers without safeguards, you could be non-compliant.

The "training" clause:

Ask: "Do you retain rights to use my data to train your base models?"

Why: You do not want your proprietary customer lists teaching a public AI how to sell to your market. Demand a "Zero Data Retention" policy.

Legacy integration strategy:

Ask: "Do we need to replace our current ERP?"

Why: If the answer is "Yes," walk away. A lean partner should know how to "wrap" your existing systems, not destroy them.

Liability & insurance:

Ask: "Do we need to replace our current ERP?"

Why: If the answer is "Yes," walk away. A lean partner should know how to "wrap" your existing systems, not destroy them.

The "drift" strategy:

Ask: "How do you handle Model Drift when our data changes over time?"

Why: AI models degrade over time as your business changes (e.g., new invoice formats). If they don't have a plan to retrain the model (MLOps), your investment will break in 6 months.

With these legal safeguards in place, the only remaining risk is financial.

The Lean approach mitigates this by replacing heavy capital expenditure with small, controlled deployments that deliver value while reducing risk.

> You don't need a massive budget; you just need a focused month.

The roadmap: Your 30-day "Lean AI" sprint

Follow this 4-week cycle to secure definitive proof of value before committing to scale.

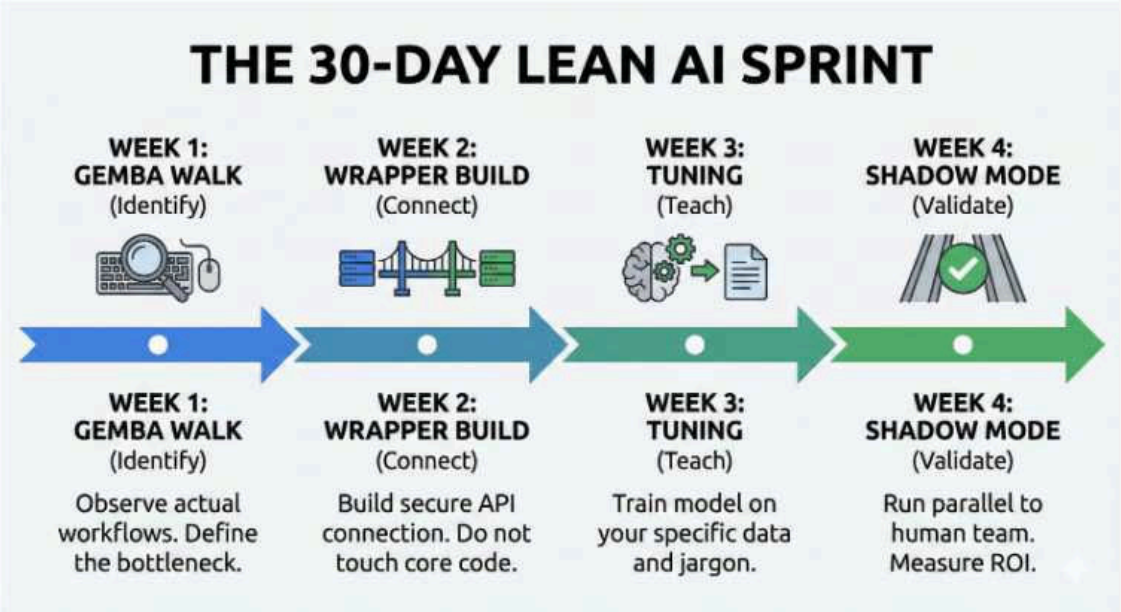

WEEK 1

The "Gemba" Walk (Identify)

Don't look at reports; look at screens. Sit with your team and find the one task where they switch between two windows to copy data. That is your target.

Goal: Define the input (e.g., PDF Invoice) and the output (e.g., Excel Row).

WEEK 2

The Wrapper (Connect)

Before building "brains," build the "bridge." Have your IT team or partner build a secure connection to that data source.

Goal: Get the data out of the silo without breaking the system.

WEEK 3

The Tuning (Teach)

Feed a sample of your data (anonymised) to the model. Teach it your specific acronyms and formats.

Goal: Get the AI to read your documents with 90%+ accuracy.

WEEK 4

Shadow Mode (Validate)

Run the new process alongside the human process. Compare the results and ask yourself: Did it save time? Was it accurate?

Goal: **Hard Data.** If it works, you have a business case. If it doesn't, you only lost 4 weeks, not 4 years..

Don't guess your week 1. Define it.

The success of your 30-day sprint depends entirely on picking the right target. By picking a process that is too complex, you will fail. And by picking one that is too simple, you won't prove value.

Need a partner to run this sprint for you?

Imobisoft specialises in executing these "Proof of Value" sprints for the UK.

Contact us