

AI from pilot to performance

A practical operating model for turning AI experiments into dependable business capability.

LE CHENE INSIGHT BRIEFING | Le Chene

The decision before the technology

Organisations rarely lack AI ideas. They lack a repeatable way to decide which ideas should become working services. A promising demonstration does not establish that an application will improve a process, fit the data available, survive exceptions or earn the confidence of the people expected to use it. The starting point is therefore an operational problem with an owner, a baseline and a measurable decision to improve.

Leaders should identify where employees spend time reconciling information, preparing recommendations, answering repeat questions or checking compliance. They should then ask what a good human decision looks like today and which part of that work AI could support. In high-consequence settings, a human remains accountable for the final decision and the organisation documents when the system must defer.

- Name the workflow and decision owner before choosing a model.
- Record the current time, quality, cost and rework baseline.
- Define the users, affected customers and non-negotiable controls.

A five-part assessment

Score candidate use cases against five questions: value, data readiness, risk, adoption and maintainability. Value means the change in a business outcome, not the number of prompts entered. Data readiness includes rights to use information, quality, access and retention. Risk covers erroneous outputs, privacy, security and unfair effects. Adoption asks whether the workflow gives staff a reason to use the tool. Maintainability asks who will monitor and improve it after launch.

The NIST AI Risk Management Framework offers a useful organising discipline through Govern, Map, Measure and Manage. Applied to a pilot, this means assigning accountability, describing the use context, testing performance and risks, then operating with controls and feedback. This is a voluntary framework, not a claim that any individual tool is certified.

- Value: Which decision becomes faster or more reliable?
- Data: Which trusted sources can the solution actually use?
- Risk: What happens when it is wrong or unavailable?
- Adoption: How will users learn, challenge and improve it?
- Ownership: Who monitors it six months after launch?

Pilot in the real workflow

A pilot should be narrow enough to learn quickly but real enough to reveal operational friction. Choose one team, one process and a defined period. Capture an accepted output standard before testing. Run representative cases, including edge cases, and compare the assisted workflow with the existing baseline. Assess not only speed but also error correction, user trust, escalation and customer impact.

Avoid treating a compelling demonstration as proof of financial value. Where work is augmented rather than removed, measure quality, capacity and service outcomes honestly. Create a register of issues and a route for staff to report misleading answers or missing data. Training should cover safe

use, checking, escalation and when not to rely on the system.

- Set an acceptance threshold and stop conditions before launch.
- Have domain experts review a sample of outputs.
- Track overrides, complaints, rework and time saved together.

Scale deliberately

Scaling requires more than a successful pilot. Standardise data access, roles, logging, support, vendor arrangements and change control. Make sure the process owner, technology team, information security, HR and affected staff agree on who owns each decision. A product owner should maintain the backlog and review whether the original value case remains valid as use grows.

The strongest AI programmes treat capability as part of implementation. Leaders need to understand investment choices and governance; managers need to redesign work; frontline teams need practice with realistic tasks. This is how a promising experiment becomes a dependable service that clients and employees can trust.

- 30 days: prioritise use cases and baseline one workflow.
- 60 days: pilot, test exceptions and review adoption.
- 90 days: decide whether to stop, refine or scale with named owners.

What this means for clients

A useful AI strategy joins business priorities, solution architecture, controls and people development. Le Chene can assemble specialists across those disciplines around a defined client problem.

Further reading

NIST, AI Risk Management Framework and Playbook

World Economic Forum, Future of Jobs Report 2025

lechene.net | enquiry@lechene.co