

Use AI where it helps. Keep accountability visible.

A practical operating baseline for AI-enabled workflows. Purpose, data, access, human oversight, third-party dependencies, documentation and measurement should be clear before automation or scope increases.

Human judgement stays responsible. Govern the workflow, not just the model.



HUMAN JUDGEMENT REMAINS RESPONSIBLE

AI can increase capacity. Accountability should remain visible.

The operating model separates what AI may prepare from what people remain accountable to decide, approve, communicate or escalate.

HUMAN

Responsible

Context, judgement, relationship, material decisions, approval and escalation.

AI

Supports

Research, classification, summarisation, prioritisation, preparation and automation where justified.

SYSTEM

Remembers

Identity, access, workflow state, history, evidence, documentation and measurement.

Not Out Labs principle

If AI does not create a meaningful contribution for the use case, it should not be forced into the workflow.

SIX CONTROLS BEFORE WIDER AUTOMATION

Purposeful data. Bounded access. Visible accountability. Reviewable change.

These controls are an operating baseline, not a certification claim and not a substitute for engagement-specific security, legal or data-processing requirements.

01 DATA MINIMISATION

Use only information necessary for the agreed business purpose.

02 SCOPED ACCESS

Limit credentials and permissions to named people, tools and systems that need them.

03 HUMAN OVERSIGHT

Define approval and escalation points for material decisions and sensitive exceptions.

04 THIRD-PARTY TRANSPARENCY

Identify material AI, hosting, automation and platform dependencies.

05 DOCUMENTATION + CHANGE

Record important workflow rules, dependencies and material changes where appropriate.


06 MEASUREMENT

Evaluate usefulness, errors, exceptions and operating impact before scope increases.

PURPOSE FIRST. MINIMAL DATA. KNOWN ACCESS.

Data handling should follow the actual workflow, not a generic promise.

Website data practices remain covered by the public Privacy Policy. Client work may require additional contractual or data-processing terms based on systems, data



IMPORTANT

Security questionnaires and data answers should be based on the actual architecture and engagement design - not broad marketing claims. Geography, providers, access, decision rights and data categories matter.

BEFORE A WORKFLOW GOES LIVE, THESE QUESTIONS SHOULD HAVE ANSWERS.

The goal is visibility, not paperwork for its own sake.

Intended use, data movement, decision rights and operating risk should be clear enough for the people accountable for the workflow.

01 What business objective is the AI-enabled workflow intended to support?

02 What data enters the workflow and which systems or providers receive it?

03 Which actions may AI prepare, classify, summarise or recommend?

04 Which decisions require human approval?

05 What errors or exceptions trigger escalation?

06 How will output quality and operating impact be measured?

07 What changes require re-review?

08 Who is accountable for access, workflow and acceptance?

TRUST IS ENGAGEMENT-SPECIFIC

Make the trust model specific to the work.

Responsible AI and data handling depend on architecture, data, jurisdictions, providers, users, access and the decisions being supported.

CLAIM DISCIPLINE

- No certification is implied by referencing a framework.
- No universal "enterprise-grade security" claim is made independent of architecture and scope.
- No assumption is made that fully autonomous AI is appropriate.
- Security and data answers should reflect the actual engagement design.

START WITH THE WORKFLOW

Purpose before platform.

Start with the business objective, intended workflow, data categories, systems, access requirements and approval points.

notoutlabs.com/trust

Human-led. Research-driven. AI-enabled. Practical. Transparent.

The role of AI should follow the real requirement, not a predetermined solution.