PROVENANT AI LABS

SINGAPORE / RESEARCH-LED ENGINEERING

Engineering, with evidence.

A focused technical partner for organisations building AI systems under real-world constraints.

THE PROPOSITION

Provenant AI Labs works at the intersection of **AI research, systems architecture and practical deployment engineering**. We help teams turn a technical uncertainty into a bounded experiment, a reproducible result and a decision about what to build next.

01

AI Architecture & Risk Review

Assess an existing system or proposed design before committing to the next build.

TYPICAL OUTPUT

Architecture findings, risk register, prioritised recommendations, evidence gaps.

02

Edge Inference Benchmark

Measure model behaviour on relevant targets and operating constraints.

TYPICAL OUTPUT

Baseline protocol, latency/resource results, comparison table, repeatable test assets.

03

Model Optimisation Feasibility

Test whether quantisation or deployment changes are worthwhile for a defined workload.

TYPICAL OUTPUT

Feasibility experiments, accuracy-performance trade-offs, go/no-go conclusion.

04

Bounded AI Proof-of-Concept

Engineer and evaluate one carefully framed capability or integration hypothesis.

TYPICAL OUTPUT

Scoped prototype, acceptance tests, implementation notes, handover record.

OUR OPERATING PRINCIPLE

Measure before scaling. Every engagement is separately scoped, has explicit assumptions and acceptance criteria, and distinguishes observed results from unproven claims.

ENGAGEMENT MODEL

A clear technical question. A useful decision at the end.

We work with the client's engineers and technical owners rather than acting as a substitute for their internal team. Commercial scope, access, work product, IP and acceptance criteria are agreed in a project-specific statement of work.

WHAT THE CLIENT RECEIVES

1 Frame
Decision to support, baseline, access boundary, measurable success criteria.

2 Evaluate
Controlled implementation or experiment, evidence capture, issue log.

3 Recommend
Technical findings, limitations, options, prioritised next steps.

4 Handover
Review session, reproducible artefacts where agreed, decision record.

WHERE WE APPLY THIS ENGINEERING

Healthcare & MedTech: privacy-conscious edge workflows, governance, model evaluation and interoperability research.
Industrial and embedded systems: constrained inference, monitoring and device-side evaluation.
Enterprise & scientific systems: AI architecture, validation methodology, data and system interoperability.

Start with a technical brief. Share the use case, constraints and target decision (not confidential datasets or regulated records). We will determine whether a scoped B2B engineering engagement is appropriate.