



BRAINPOOL.AI

WEBINAR • 16 JUNE 2026 • 1-2PM

Own Your Intelligence

Why 95% of AI projects fail, and what the environmental sector can do differently.


A working session on building AI you own, not AI you rent.

BEFORE WE BEGIN


Your feedback on AI to date

344 of you registered. Most already use a generic tool, and only one in four is satisfied with it.

What you're using today



Happy with your current tools?



“

The outputs sound so generic - like corporate speak. I have to completely rewrite everything.

Consultant

“

It lies and makes stuff up. Copilot hallucinates constantly.

Project lead

“

They fall over on CSMs - they miss the nuances of geochemical and migratory relationships.

Technical director

MIT SLOAN MGMT REVIEW

95%


of AI projects fail to reach production.

Why AI fails in your sector

The same five patterns kill AI projects - and you told us all five are happening here.

01	Accuracy collapse	AI confidently fabricates regulatory thresholds, CSM pathways and EIA citations. More time checking than the task took – net negative.
02	Generic tools, specialist work	Copilot wasn't trained on Phase 1 protocols, DEFRA metric logic or your firm's standards. Plausible-sounding nonsense only experts catch.
03	The trust gap	When you can't verify a result you check everything manually. Without auditability, adoption stalls at pilot stage.
04	Vendor lock	Most of this room is restricted to Copilot only. Locked-down tools with zero domain training are the key barrier.
05	Experiments that never ship	Pilots get interesting results, then die. No ownership means no ability to improve, scale or maintain the model.

THE PRODUCTIVITY PARADOX



A single query is cheap - until it starts to “think”

0.24_{wh}

Median Gemini text prompt — the only methodologically disclosed figure (Google, 2025).

4–40_{wh}


A reasoning / extended-thinking query — one to two orders of magnitude higher.

>200×


Spread between the lightest and heaviest models — model & mode choice dominates.

Source: [Brainpool AI — The Energy & Greenhouse Gas Footprint of LLM Inference](#)

Energy tracks tokens generated, not model size. The everyday queries that dominate usage are tiny; the cost spikes only when a model is asked to reason.



Each query gets cheaper — yet total energy roughly doubles by 2030

**33× / 280×**

Per-prompt energy fell 33× in a year; per-token cost fell 280× in two (Google).

×2

Data-centre electricity ~415 → ~945 TWh by 2030 — 4× the pace of overall demand (IEA).

~75%

Share of AI energy that is inference, not training, by 2030. The recurring bill is what scales.

Jevons

Cheaper queries beget more queries. Efficiency isn't the brake — it's the accelerator.

Containing AI's footprint is a grid-decarbonisation and procurement problem — not one efficiency alone can solve.

Sources: IEA Energy & AI (2025); Google / MIT Tech Review (2025); AIMultiple. Projections are modelled scenarios with wide ranges.

CUSTOM & FINE-TUNED LLMs

How to cut energy per query

Decode dominates. Output generation is ~16–26× the energy of input — control tokens and reasoning first.

1 Serving layer

No model change

Continuous batching (vLLM)
up to 23× throughput

Prompt / prefix caching
up to 90% input cut

Semantic caching
skips 30–60% of calls

Cap output & reasoning
~70% on reasoning

2 Routing & FP8

Low risk

Model routing / cascades
85% cost, 95% quality

Quantise to FP8
~11–30% energy/token

INT4-AWQ (VRAM-bound)
~2.7× speed

3 Model changes

Amortise over volume

Right-size / distill
40% smaller, 97% quality

Prefer MoE base
5.5% params active

Fine-tune to drop prompts
past break-even only

4 Infrastructure

Deploy & site

Newer accelerators (H200/B200)
35–45% energy/token

Maximise utilisation
idle GPUs ~10× waste

Low-carbon grid region
cuts CO₂, not energy

Biggest win: serving-layer changes need no retraining and stack to ~8–20× combined energy/query reduction. Fine-tuning helps mainly by enabling a smaller model or shorter prompts — not by itself.

THE POWER OF CUSTOM AI • 01

Regulatory Compliance Intelligence

A guidance-currency engine for permits and assessments.



WHAT IT DOES

AI reviews a draft report, permit application or assessment against the live body of regulatory guidance - Environment Agency, EPA, SEPA, Natural England and the relevant statutory instruments - and flags where the work leans on superseded thresholds, withdrawn guidance or the wrong methodology. It cites the specific clause it checks against, so a consultant verifies rather than trusts blindly.

“Mostly, they can be factually inaccurate.”

Webinar registrant

ACCURACY
Y

WHY IT HAS TO BE CUSTOM



Guidance never stops changing

Off-the-shelf models freeze the day they ship and silently drift out of date. Fine-tuning and model training ground every answer in a continuously-maintained, version-controlled corpus of live EA / EPA / SEPA sources.



Traceable, not plausible

Answers retrieve from an owned knowledge base - not the model's memory - so each one points back to a current clause rather than a confident hallucination.



Liability stays private

Regulatory positions carry legal liability, so the whole pipeline is owned and deployed privately - never routed through a public API.



THE POWER OF CUSTOM AI • 02

Technical Report Generation

Grounded drafting for contaminated-land, EIA and monitoring reports.



WHAT IT DOES

AI drafts the repetitive, high-volume sections - site setting, conceptual site models, monitoring summaries, data tables, methodology - directly from your own inputs: borehole logs, lab results, GIS layers and prior reports. It writes in your house methodology and tone, and every quantitative statement traces back to a source record.

“Improving all the time — but lots of ghost data; we need to sense-check everything.”

Webinar registrant

TRUST

WHY IT HAS TO BE CUSTOM



Kill the ghost data

The model is constrained to your validated data - it generates from real lab values and logs, not statistical guesswork. No invented figures in a liability-bearing report.



Your methodology is your IP

Report library and house style are your competitive edge. Training a generic SaaS tool on them hands your differentiator to the vendor — and every competitor on it.



Own the model and the data

Both stay inside your deployment. Here, “own your IP” isn't a slogan — it's the only viable architecture.



THE POWER OF CUSTOM AI • 03

Bid Scoping & Quotation Engine

Automated scoping and pricing for tenders — the fastest path to measurable value.



WHAT IT DOES

AI system ingests an incoming RFP, breaks it into a scope of work, estimates effort and cost against your historical project economics, identifies risks and exclusions, and drafts the proposal. It's Brainpool's proven entry point — and the most immediately measurable in time saved.

“How does a time-and-materials business capture the value of what we provide as we become more efficient?”

Webinar registrant

STRATEGY

WHY IT HAS TO BE CUSTOM



The intelligence is proprietary

Accurate scoping depends on your cost base, win/loss history, day rates and what similar jobs actually took to deliver — none of which exists in a public model's training.




Escape the efficiency trap

Bill by the hour, get faster, earn less. A quotation engine shifts the firm toward value-based scoping — but only when trained on your own economics.



Confidential by necessity


That economic data can never live in a shared, generic tool. It has to be private, and it has to be owned.



THE STAKES

Sometimes you only have one chance to get it right.

Environmental decisions are often irreversible and legally binding. A confidently-wrong answer isn't an inconvenience - it's liability. A failed permit, or a remediation that can't be undone.




YOUR AGENTIC AI PLATFORM



How to achieve production-grade accuracy with end-to-end customisation.

A continuous loop — every cycle curates data that distils a smaller, cheaper, more accurate model.



WHY OWNERSHIP COMPOUNDS


Four moats a custom model builds



01

Data moat

Trained on your projects, standards and approved deliverables — institutional knowledge no competitor can ever see.



02

IP ownership

You own every improvement. The model is a permanent asset on your balance sheet — not a subscription you rent.



03

Security & privacy

Deployed in your environment. Client process data and crown-jewel project files never leave the building.



04

Compounding edge

Every project makes it sharper. The longer it runs, the harder it becomes for anyone else to replicate.

TWO STRATEGIES

Rent the model, or own it

Two fundamentally different AI strategies. Only one builds lasting value.

Rented

generic, off-the-shelf

- ✗ Trained on generic internet data
- ✗ The vendor owns every model improvement
- ✗ Data security is unclear
- ✗ Switch vendors created technical debt
- ✗ Every competitor has the same tool
- ✗ No competitive differentiation

Owned


custom · Brainpool

- ✓ Trained on your projects & standards
- ✓ You own all improvements and IP
- ✓ Data stays in your environment
- ✓ A permanent asset that grows with you
- ✓ Competitors cannot replicate it
- ✓ Becomes a measurable firm asset

THE JOURNEY

Start generating value from AI in 3–6 months

A tested, step-by-step process designed to de-risk your investment and accelerate time to value.



CROSS THE GENAI DIVIDE

Thank you



BRAINPOOL.AI

| kasia@brainpool.ai

peter@brainpool.ai