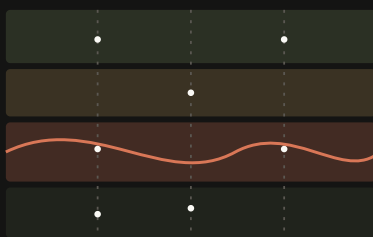


VERTICALLY-INTEGRATED AGENTIC AI


# A company that solves real problems **end to end.**

One rare capability – agents whose confidence is **calibrated** and whose autonomy is **earned** – put to work end to end, one vertical at a time.


Mining is our first home, taken A to Z. Telecom is the first proof the same engine transfers. It is built to go further.



HETEROGENEOUS SUBSURFACE DATA



QUANTIFIED UNCERTAINTY




DRILL-READY DECISIONS

# At a glance

---

A two-founder, vertically-integrated agentic-AI company – registered in Singapore, with a working platform, five published public-data case studies, and engagements across three countries, all in roughly three months.



Idea → a registered, signed Singapore company in roughly three months.

## BUILT FAST, ON PURPOSE

In roughly three months: **incorporated in Singapore**, built the platform, published **five public-data case studies** (live on [tolun.ai](https://tolun.ai)), and opened engagements across **three countries**. The pace is the point – **agentic-AI speed, paired with deep domain knowledge and a mission.**

# 01

MISSION · PLATFORM

## Why we exist, and what we are

---

We exist to take the weight off human work – to end the rounds of wasted effort, costly wrong calls, and broken communication that grind people down.

Our aim is to put hard-won knowledge within easy reach. The discipline that keeps our AI honest – it refuses to fake a result, shows its uncertainty, and a human always signs – is that mission written in engineering.

### WHAT THE PLATFORM SETS OUT TO DO



#### Solve mining, A→Z

The whole arc, not exploration alone: find the deposit, govern the reserve, plan, operate, and close – with a Competent Person signing every call.



#### Lighten the load on people

Agentic AI that carries coordination across sites, disciplines and time zones – not to replace anyone, but to give teams back ease and confidence.



#### One engine, many domains

The same calibrated-autonomy engine runs on any data-and-decision-heavy field. Mining and telecom are the first two deployments – not the limit.

### POSITIONING

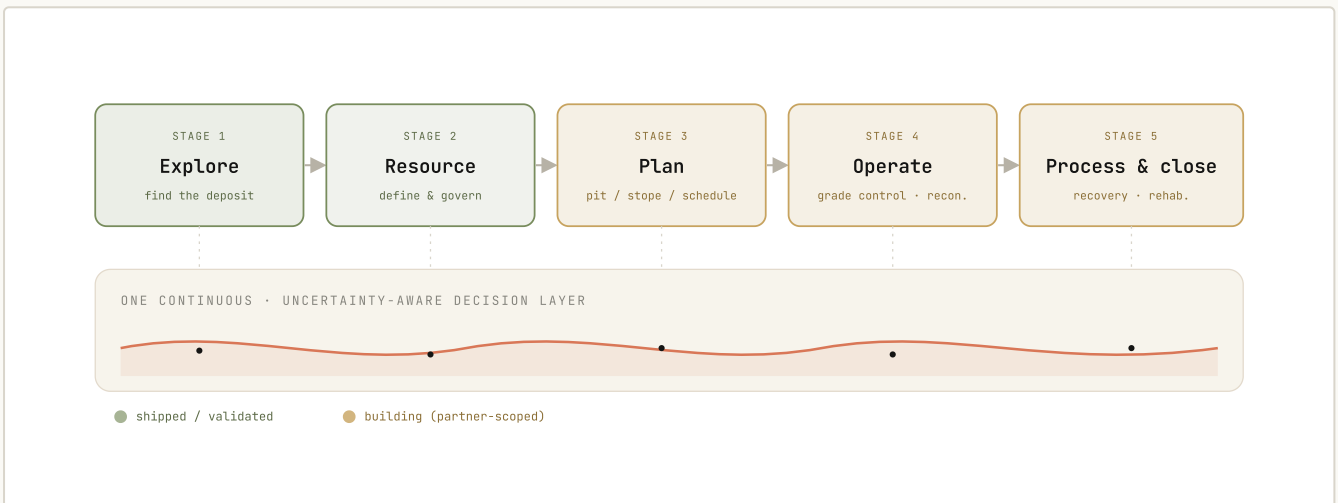
**The engine is the product; verticals are deployments.** A vertically-integrated agentic-AI platform that solves mining first – not a mining-software company, and not a single-technique tool. Domain-agnostic by design.

# 02

A → Z

## The mining lifecycle, end to end

Every incumbent tool is single-stage. Our platform spans the whole orebody lifecycle – one continuous, uncertainty-aware decision layer from first survey to mine closure.



One engine across five stages – the same data answering a different question at each step.

2 / 5

STAGES SHIPPED & VALIDATED

3 / 5

STAGES BUILDING · PARTNER-SCOPED

1

CONTINUOUS DECISION LAYER

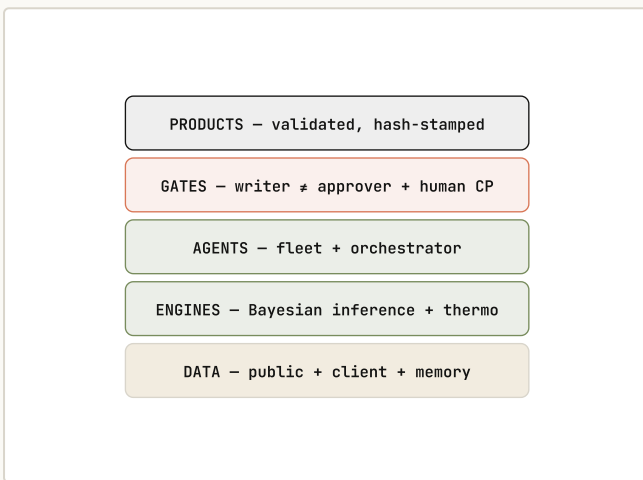
**CONTINUITY** The advantage isn't any single stage – it's that **the same uncertainty travels across all five**, so a decision at closure can be traced back to the assay that started it.

# 03

BUILT VS ROADMAP

## The platform – five honest layers

One stack, bottom to top: data feeds the engines, engines feed the agents, gates hold the agents honest, and only validated, hash-stamped products leave the building.



Bottom to top: data → engines → agents → gates → products.

CAPABILITY	STATUS
Geophysics engine	BUILT
Geochem modal engine	BUILT
Thermodynamics / property	BUILT
Agent fleet + gates	BUILT
Learned prospectivity	PROTOTYPE
Joint geochem + geophysics	ROADMAP

### THE HONEST LINE

Four engines are built and in the repository. Learned prospectivity is a prototype. The **joint** model that fuses chemistry and physics is the roadmap centrepiece – three pieces exist, the fusion is the work.

4

ENGINES BUILT & IN REPO

1

PROTOTYPE · LEARNED PROSPECTIVITY

1

ROADMAP · JOINT MODEL

# 04

HONESTY BY CONSTRUCTION

## What makes it different

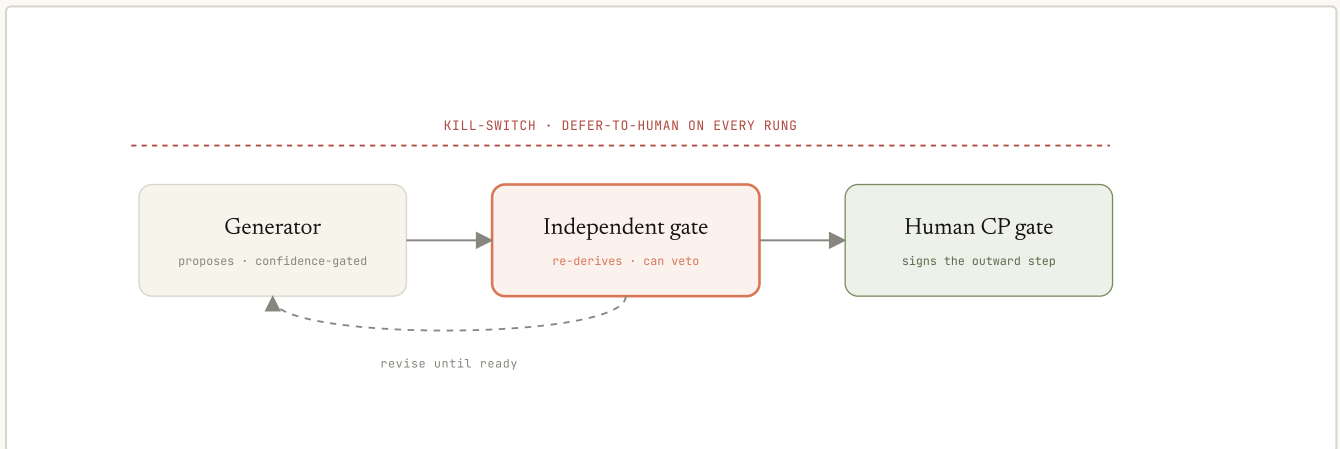
### EVERY ANSWER CARRIES ITS UNCERTAINTY

We never ship a bare number. Each result is a full posterior – the answer **and** the band around it – and the platform states plainly what the data **cannot** resolve.

On a real lithium deposit, our chemistry-to-mineralogy model recovered the economic budget, then named the mineral split the chemistry could not resolve – rather than inventing one.

### HONEST-AGENTIC LOOPS: WRITER ≠ APPROVER

The agent that writes never approves. An independent gate checks it; the writer revises; a human Competent Person signs anything irreversible or outward-facing. The writer and the approver, by design, are never the same.



The loop is not cosmetic – it has caught real errors before they shipped.

### IN ONE WEEK, THE LOOP CAUGHT – BEFORE A CLIENT SAW THEM

#### ● CAUGHT 01

**False geochemistry result**

#### ● CAUGHT 02

**A 10× mass error**

#### ● CAUGHT 03




**An over-claimed correlation**

# 05

3 COUNTRIES · STATED BY TRUE STATUS

## Traction

Real traction for a two-founder firm – presented as engagements, not a roster of paying clients we don't yet have.

<b>Australia</b> LITHIUM  DELIVERED	<b>Saudi Arabia</b> PROSPECTIVITY  ACTIVE	<b>Indonesia</b> COAL  CHANNEL PILOT	
<ul style="list-style-type: none"><li>● AUSTRALIA · LITHIUM · DELIVERED</li></ul> <b>Chemistry → mineralogy</b> <p>Routine drill assays turned into per-interval modal mineralogy with quantified uncertainty – no new sampling. Blind-tested against TIMA: the lithium-silicate budget recovered to within ~9% at the deposit mean.</p>	<ul style="list-style-type: none"><li>● SAUDI ARABIA · PROSPECTIVITY · ACTIVE</li></ul> <b>Stage-1 target generation</b> <p>A live, multi-commodity Stage-1 assessment over the Jeddah / Bahrah ground from public and government data – evidence-ranked corridors, gated by an independent specialist vet.</p>	<ul style="list-style-type: none"><li>● INDONESIA · COAL · CHANNEL PARTNER (PILOT)</li></ul> <b>Wireline-density voxel + QC</b> <p>Stage-B/C wireline-density modelling for shallow coal, surface-gravity as an independent QC layer. A partner pilot, not a signed paying client – framed that way.</p>	<ul style="list-style-type: none"><li>● AUSTRALIA · PUBLIC DATA · PUBLISHED</li></ul> <b>5 independent validations</b> <p>Blind public-data case studies on tolun.ai: <b>Cadia-Ridgeway &amp; Northparkes</b> (porphyry Cu-Au), <b>Winu</b> (Cu), <b>Pilgangoora</b> (Li), <b>Hillgrove</b> (Sb).</p>

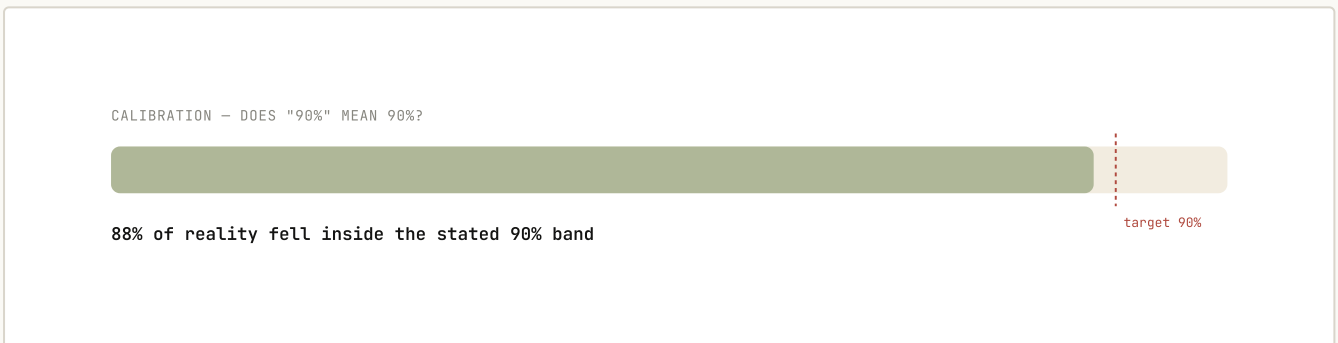
# 06

DOMAIN-AGNOSTIC · REAL POC

## Proof the engine is domain-agnostic

To test whether the platform's honesty is specific to rock, we took it out of geoscience entirely and pointed it at telecom network demand.

The scarce thing – confidence that is honest, autonomy that is **earned** – isn't specific to rock.



Fatimah's real run on open Milan mobile-traffic data (ODbL). The load-bearing result is the calibration.

**88%**

INSIDE THE STATED 90% BAND

**+29%**

SKILL VS PERSISTENCE

**268**

MAE · REAL RUN

### WHY THIS IS PROOF, NOT A PIVOT

We didn't rebuild the company for telecom. We carried the same calibrated-inference + defer-to-human machinery into a foreign domain and it behaved correctly on independent public data – the context for a bounded, certified subcontract role in the telecom-automation ecosystem.

# 07

SCIENTISTS & ENGINEERS

## Two founders – the two halves the problem needs

One builds the machine that thinks; one is the scientific authority that signs. Neither half solves the problem alone.



### Dr Alim Polat

CO-FOUNDER · PRINCIPAL AI / LLM ENGINEER


Builds the platform – inference engines, agent fleet, and the honest-agent loops. A principal engineer with over a decade of industry experience in large-scale AI systems.



### Dr Fatimah Abdulghafur

CO-FOUNDER · PRINCIPAL GEOSCIENTIST · COMPETENT PERSON

Carries the domain authority – a PhD geoscientist and Competent Person with real industry experience, the scientific gate of record on every deliverable, and author of the telecom PoC.



Writer and approver, by design, are never the same person.

# 08

THE ASK


## Where this goes

Tolun AI is a two-founder, vertically-integrated agentic-AI company – now a **registered Singapore company (Tolun AI Pte. Ltd.)**, with a platform, five published public-data case studies, and engagements across three countries, all in roughly three months. Its moat is uncertainty-quantified inference wrapped in honest-agentic loops.

### THE ASK

A conversation, and a bounded first engagement where the calibration and defer-to-human behaviour proven on public data is demonstrated on a sanctioned, non-critical feed of yours – small enough to start now, honest enough to walk away from cleanly.

### ONE TRUST LAYER, MANY DOMAINS



# TOLUN · AI

COMPANY

Tolun AI Pte. Ltd.

REGISTERED

Singapore · UEN 202628335N

WEB

tolun.ai

FOUNDER

Dr Alim Polat

Principal AI / LLM Engineer

FOUNDER

Dr Fatimah Abdulghafur

Principal Geoscientist · Competent Person

Engagements are stated by true status. The only quantified performance figures shown are from a real, independently-sourced public-data run.