MINOTAUR

Fund Introduction

Minotaur Capital • Global Opportunities Fund • August 2026

WHERE WE'RE HEADED

From a philosophy, to a machine that runs it.

1 **Who we are:** the record so far — then one idea, shown on two real positions

2 **How it evolved:** from one lucky article to a repeatable system

3 **The AI team today:** twenty-seven analysts, and how they work

4 **Why you can trust it:** fact-checks, benchmarks, version control

5 **Risk & the portfolio:** how it all comes together

WHO WE ARE

An AI-first fund, built from scratch.

- Founded November 2023; launched the Global Opportunities Fund in May 2024.
- Global equities: every country, sector and size; no style box.
- Two founders — Thomas Rice and Armina Rosenberg — 35+ years combined funds-management experience.
- We run our own platform, Taurient. The same software does the research and the portfolio.
- The rare bit: at Minotaur the investor and the engineer are the same people, so the machine encodes *our* thinking, not a generic tool bolted onto an old process.

THE HUMANS BEHIND IT

Two founders, one unusual combination.

Armina Rosenberg

CO-FOUNDER • PORTFOLIO MANAGER

- 8 years sell-side research at JPMorgan
- 2 years at a prominent family office
- 5 years running global equities for one of the Atlassian founders

Thomas Rice

CO-FOUNDER • PORTFOLIO MANAGER

- 11 years rising through the ranks at PM Capital
- 9 years at Perpetual, 6 running the Global Innovation Fund
- Coding since childhood — built Taurient, the system in this deck

A true team: diverse in background, complementary in skillset — the investor and the engineer under one roof.

TRACK RECORD

Two years in.

PERIOD	FUND	MSCI ACWI	ALPHA
3 months	+8.2%	+6.9%	**+1.3%**
1 year	+10.8%	+11.9%	−1.1%
2 years p.a.	+18.8%	+14.7%	**+4.1%**
Since inception p.a.	+19.0%	+16.1%	**+2.9%**
Since inception	**+47.3%**	**+39.5%**	**+7.8%**

As at 31 July 2026. Returns inclusive of fees, in AUD. Inception 9 May 2024. Past performance is not a reliable indicator of future performance.

TRACK RECORD · ATTRIBUTION

Where the returns came from.

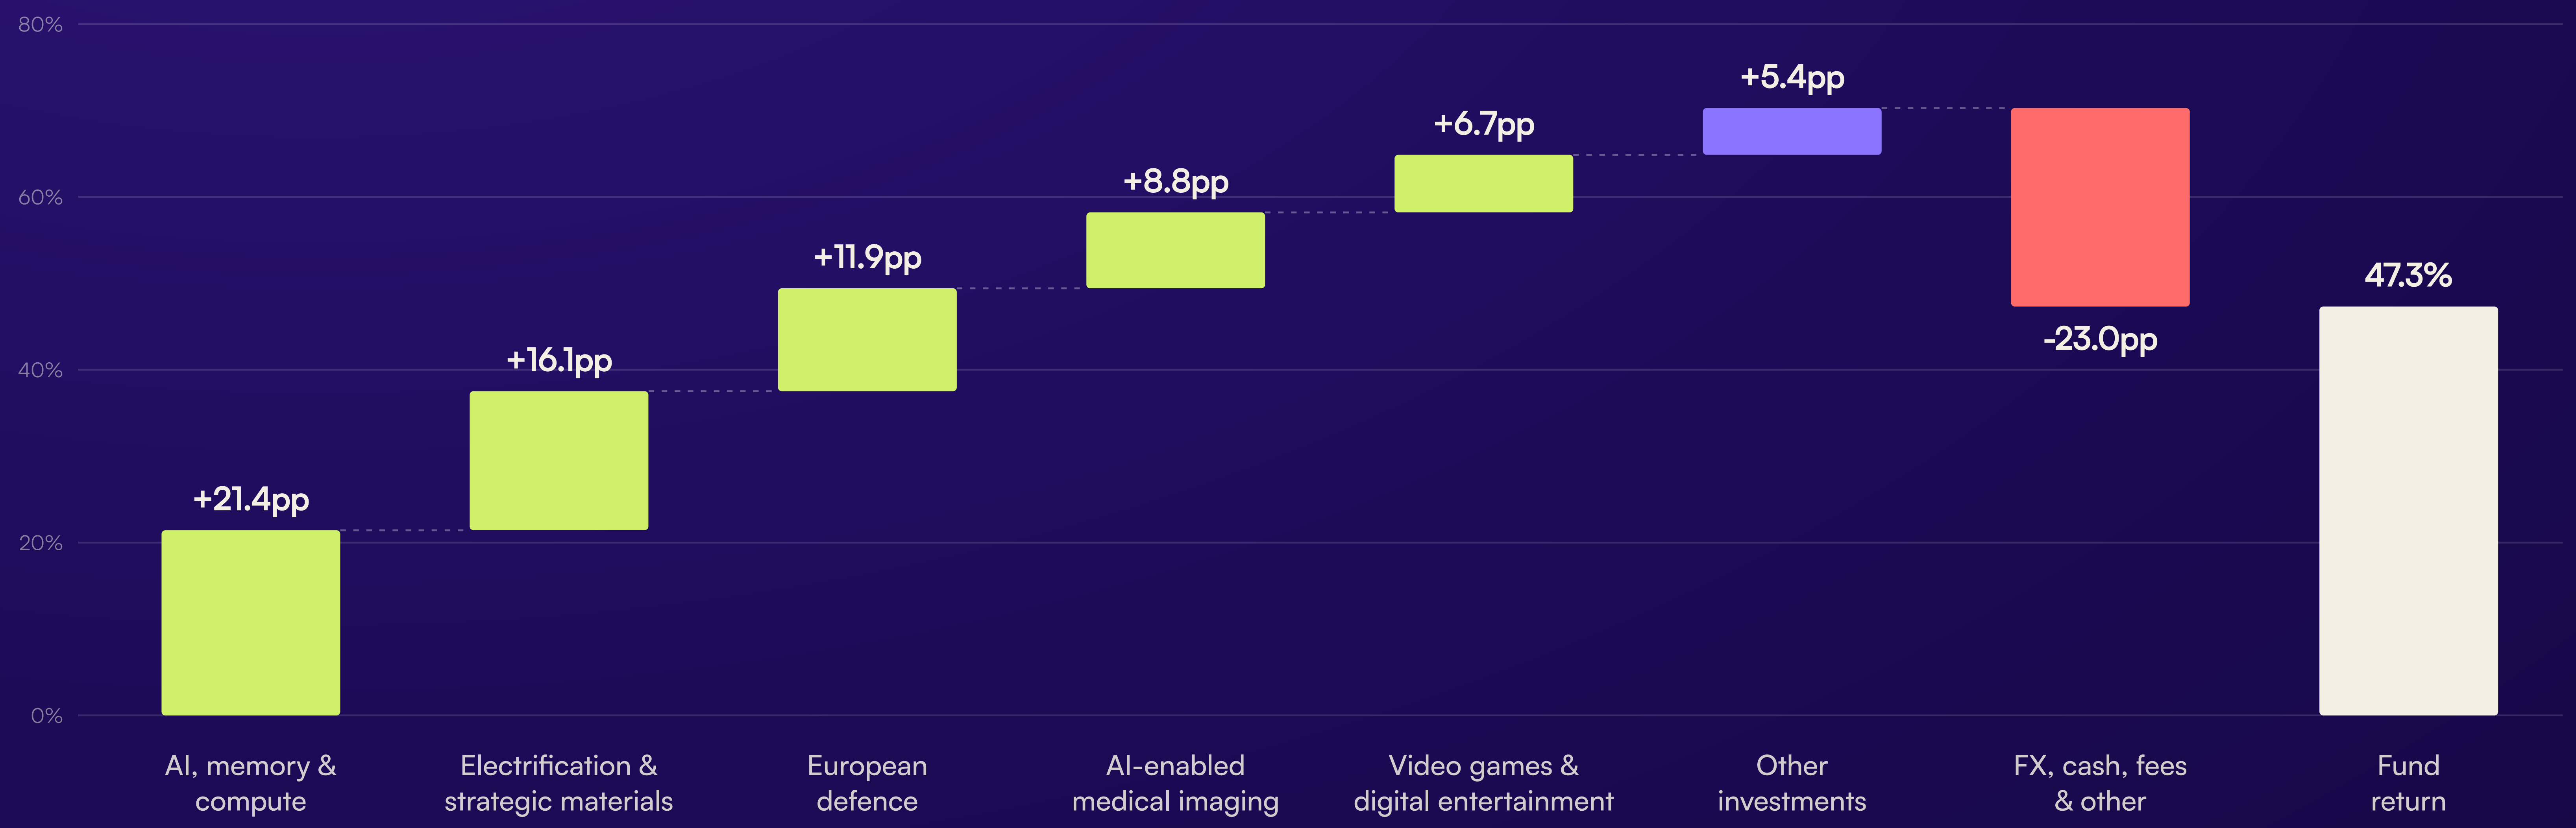

Estimated contribution, 9 May 2024 inception to 31 July 2026. Themes are judgement-based groupings of security-level contribution. Reconciliation includes foreign exchange, cash, fees, tax and other calculation differences. Past performance is not a reliable indicator of future performance.

THE FIRST QUESTION EVERYONE ASKS

Growth fund? Value fund? Tech fund? The label is the problem.

A style label is a pre-packaged set of mental shortcuts — it quietly commits you to the wrong questions before you've looked. So we don't pick a box. We ask the same four questions of every company on earth:

1. How does this business **make money**?
2. What will it **earn over time**?
3. How **certain** is that?
4. What is it **worth**?

One discipline, any sector, any region, any "style". The fund fits none of the boxes — deliberately.

THE ONE IDEA BEHIND EVERYTHING

We hunt the market's wrong model.

A business gets mispriced when the market is running the **wrong model** of it. We build the right one, from first principles.

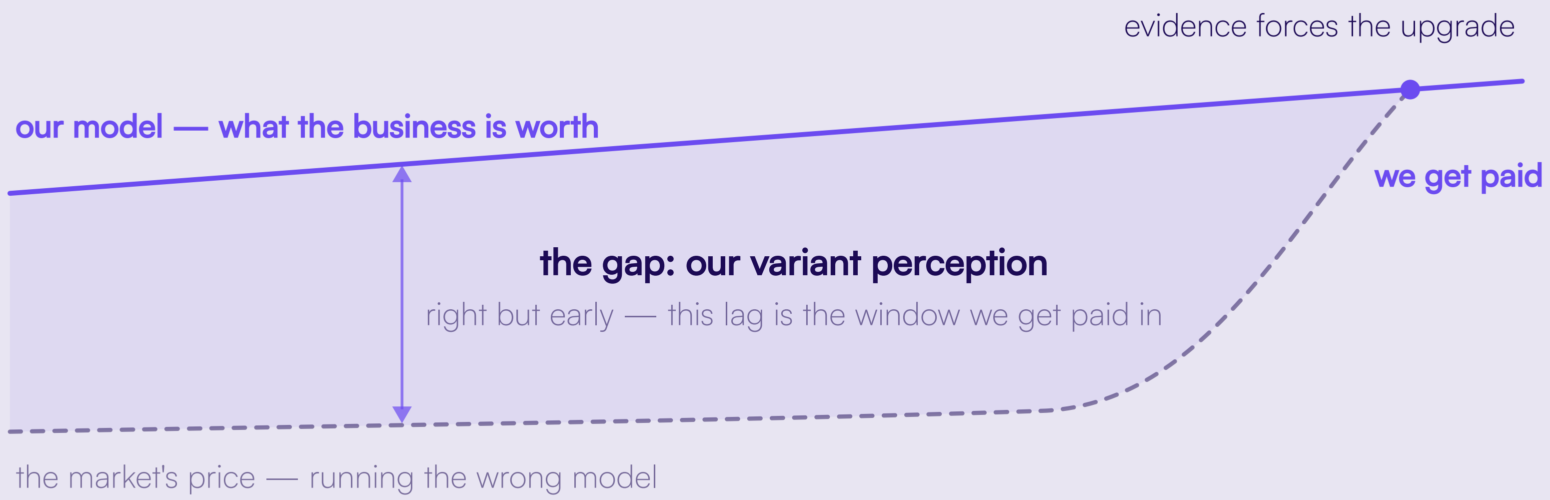

We never hold a **consensus view at a discount**: every position embeds a **variant perception** — a specific, researched reason the market is wrong.

> The market has run the wrong model of **Google** twice in our careers: at the IPO — "an internet advertising company that should trade at a discount to Yahoo" — and again in the 2025 AI-search panic. **We got paid both times.**

FIVE WAYS A MODEL CAN BE WRONG

The same opportunity, five shapes.

The market is mostly efficient — so the real question is where you can find the inefficiencies that drive alpha on a consistent, repeatable basis. Our answer: the same five shapes, over and over.

1 **Mis-categorisation**: filed in the wrong box, so the people who could value it aren't the ones who own it. *Axon, dismissed as a stun-gun maker while it built a body-camera subscription.*

2 **Controversy**: the market is loudly split and the wrong side leads the price. *Meta at $90 in late 2022, priced as though the metaverse would consume it.*

3 **Lossy heuristic**: valued on the wrong metric. *CD Projekt looks expensive on this year's earnings, cheap on a full development cycle.*

4 **Horizon myopia**: the market won't price what's more than a year or two out. *SEEK's decade of pricing power.*

5 **Right model, wrong numbers**: the framework's fine, the inputs aren't. *Memory is cyclical; this cycle just runs longer than consensus thinks.*

"It's already run" is not a valuation.

THE CHANGE · FEBRUARY 2025

A real shift in political will

The Munich Security Conference signalled a genuine move toward European defence autonomy — a material change to the industry's future, not the familiar "Europe should spend more" refrain. We had never focused on the sector. We didn't need to have: change is the trigger, not familiarity.

THE WORK · A 2030 VIEW

The chart doesn't set the price

The names had already risen sharply — exactly where most investors decide they've "missed it". But a past move says nothing about value. We worked through what Europe could realistically spend by 2030, built revenue, margins and earnings bottom-up, and concluded the businesses were still fundamentally cheap.

It became a meaningful position — Rheinmetall, BAE, Dassault and RENK among the names — and one of the key drivers of returns since inception. The shape: horizon myopia, a market refusing to price a multi-year future. The cleanest single win: RENK, a tracked-vehicle drivetrain monopolist the market had filed as just another cyclical — closed a year later at nearly double.

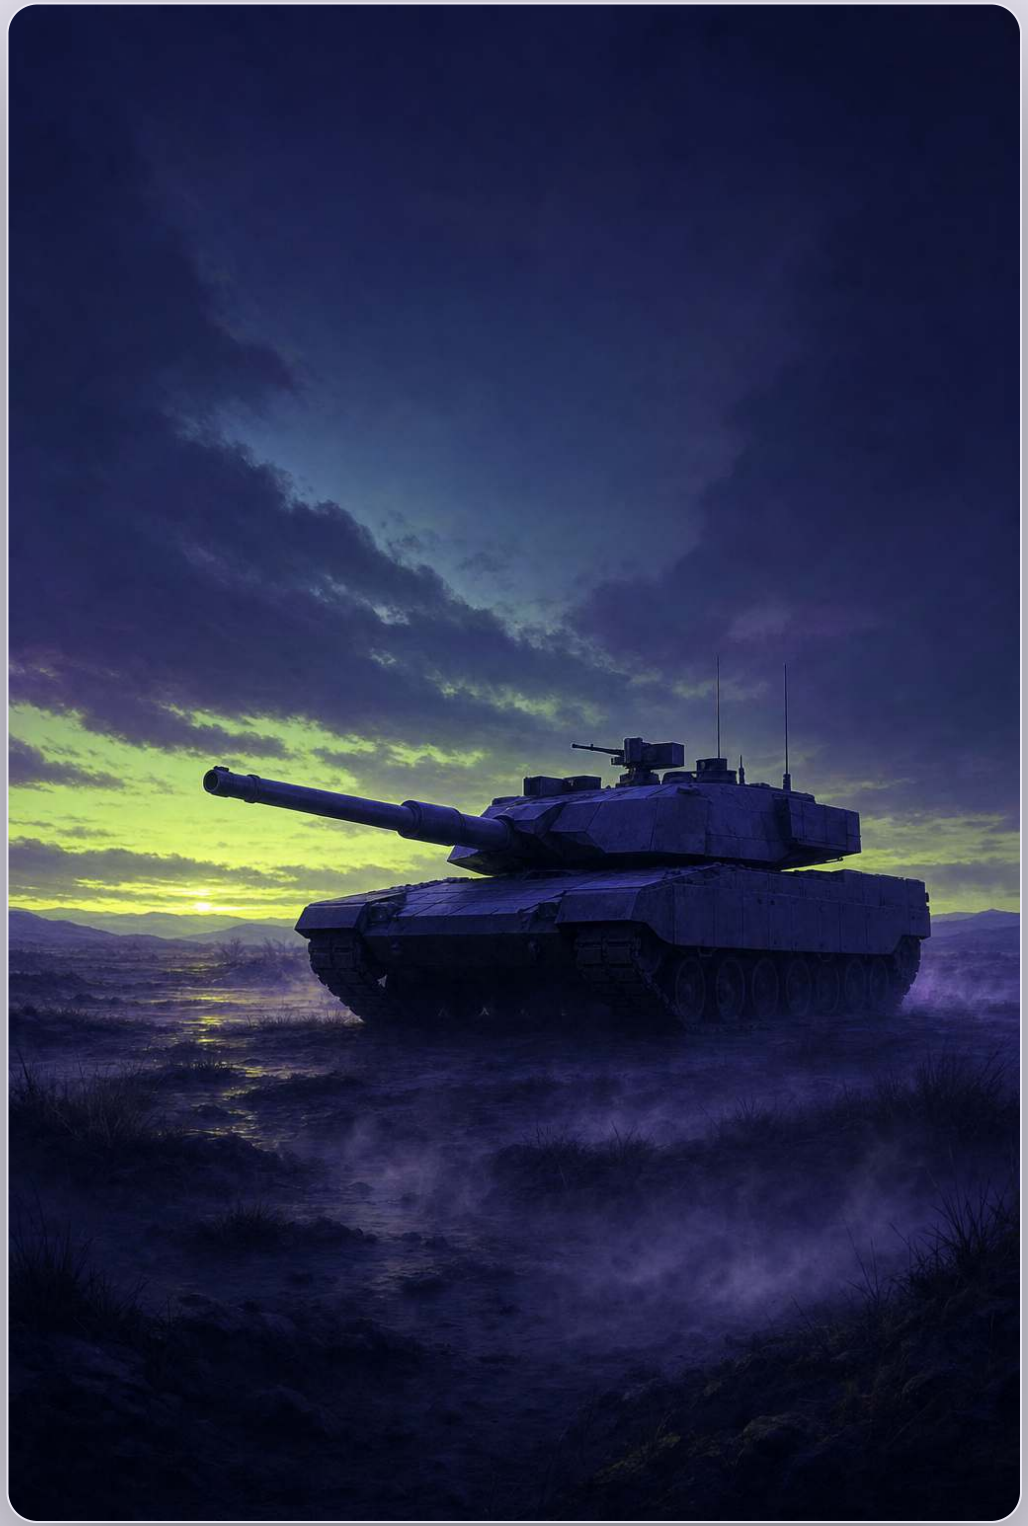

Expensive on the lazy metric. Cheap on the right one.

THE MARKET'S METRIC

Forward P/E says "expensive"

A hit-driven game developer looks most expensive on next year's earnings precisely when it's investing hardest in the next hit — the launch year is still two or three years out. A one-year multiple is a lossy compression of the real economics.

OUR MODEL

Cheap across the full cycle

A game ships every few years, so earnings peak and trough around each launch. Averaged over the whole cycle, the economics are among the best in global gaming — and on cycle earnings, today's price is cheap. It's a multi-year *Witcher IV* view that a 12-month consensus model structurally cannot hold.

The shape: a lossy heuristic, plus horizon myopia. A live position today, underwritten on a multi-year view.

The Witcher — CD Projekt's flagship franchise (illustration)

WHY OUR PORTFOLIO LOOKS LIKE NOTHING YOU'VE SEEN

One discipline. Five shapes.

SK Hynix
Korea • memory

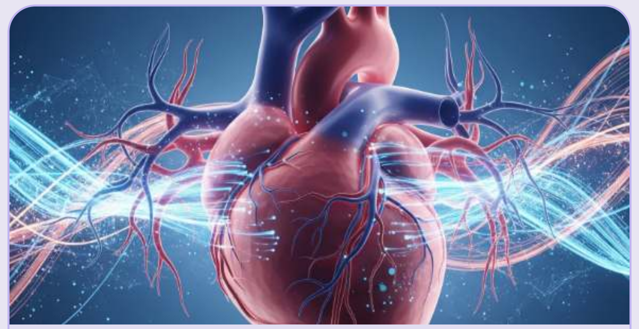

Artrya
Australia • medtech

UniCredit
Italy • banking

CD Projekt
Poland • games

Eli Lilly
US • pharma

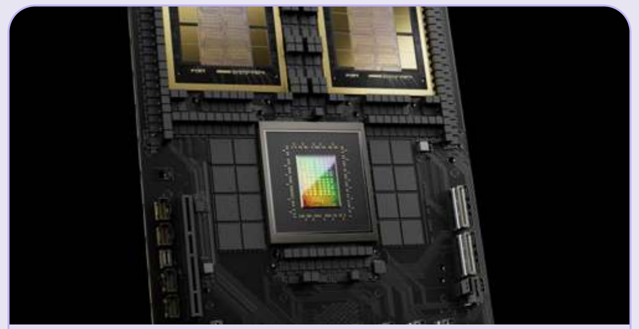

NVIDIA
US • AI compute

Crocs
US • footwear

HCA Healthcare
US • hospitals

NextEra
US • power

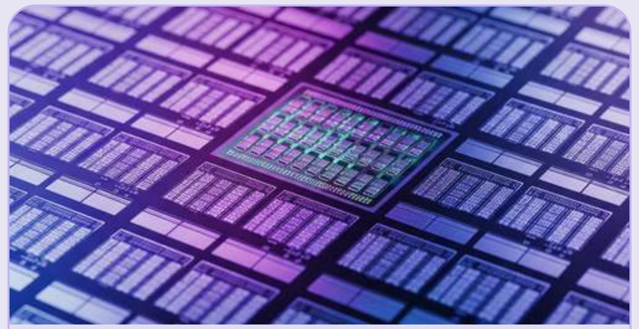

Micron
US • memory

The top ten — **nothing on a style sheet connects them**. Underneath, one thing does: in every single one, a researched case that **the market's model is wrong**.

A philosophy is only worth what you can repeat — so we turned it into a process, and wrote the process into software.

Top 10 holdings as at 31 July 2026.

THE CRAFT — DONE THIS WAY FOR FIFTY YEARS

The job, and the problem with it.

The universe those ideas come from — 58,000+ listed companies (each dot ≈ 20)

The edge was never judgement — it was **breadth**. So our approach has been to ask: **what can AI replace, and where can AI enhance?**

THE OBVIOUS QUESTION

Does AI know how to invest?

No.

A brilliant brain in a jar. It doesn't replace judgement — it lets two people's judgement **scale**. What matters is what you wire around it:

1. **Philosophy**
   A theory of how you outperform.
2. **Process**
   The theory, written down and repeatable.
3. **Automation**
   Decide what AI runs — and what stays ours.

WHAT TWO AND A HALF YEARS OF THAT BUILT

A machine that researches the way we do.

It reads the world's news, triages ideas with our philosophy, and works around the clock. What it grew into is the best part of the story — but it starts with **one lucky article**, back in 2017.

HOW IT EVOLVED · THE ORIGIN STORY

Our best-ever investment, we found by accident.

2017 — One evening, a single article: Taser was renaming itself **Axon** and giving body cameras away as "a network play".

the market — Still owned it as a **stun-gun maker** — and missed the recurring body-camera subscription business underneath.

the result — It ran **20—30×**.

the lesson — It would have been easy to miss. It depended on **luck**.

→ So when we built Minotaur, the question was: how do we turn that one piece of luck into a repeatable system?

TURNING SERENDIPITY INTO A PROCESS

Read everything. Triage fast. Go deep where it's real.

~40,000 articles / week
FLAG MATERIAL CHANGE

→

Idea Triage
TWO ANALYSTS VOTE

→

Snapshot
FAST FIRST READ

→

Initiation
FULL DEEP DIVE

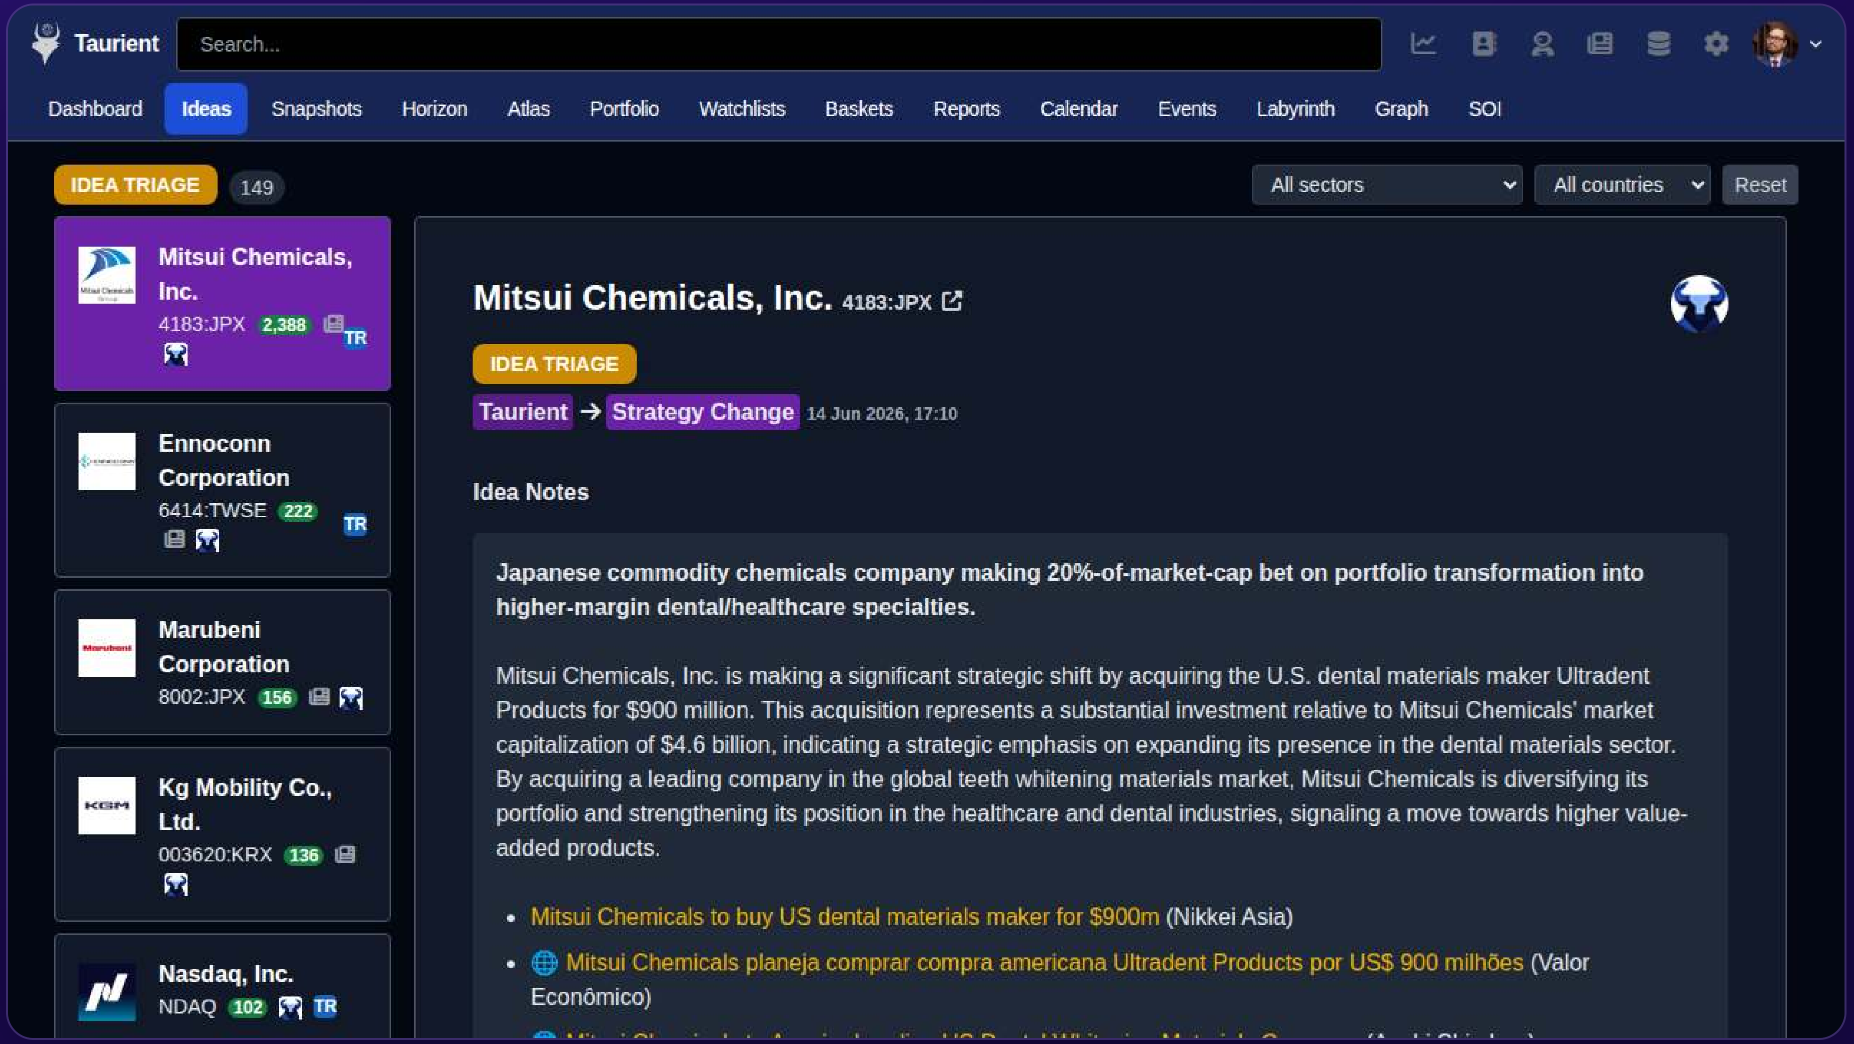

We scan for companies going through **change** (a strategy shift, new management, a new business line) because change is where a business's three-to-five-year future is genuinely up for debate, and where a fundamental investor adds the most value.

HOW IT COMPOUNDED

One capability at a time.

Nov 2023	**News scanning** and **Idea Triage** and **Snapshots**.
Jul 2024	First automated **earnings summaries**.
Jan 2025	**ChatGPT Deep Research launches.**
Aug 2025	**Research moves into git** — so AI can edit research directly, the same way we edit code.
Sep 2025	**Labyrinth Agents** — drop an `ai.md` in a stock's folder and an agent watches the news daily, **updating the research itself**.
Nov 2025	First **automated initiation**.
Jan 2026	**Kairos** — the first named agent (scheduling).
Feb 2026	First **AI-built financial models**.
Apr 2026	**Talos**, the first AI analyst — with **daily autonomous monitoring** of his patch.
Aug 2026	**32 named agents** across research and operations.

No single leap — each capability made the next one possible. The newest layer, **the analysts**, is where it compounds.

THE AI TEAM TODAY

Not one big AI.
A team of specialists.

meet the analysts →

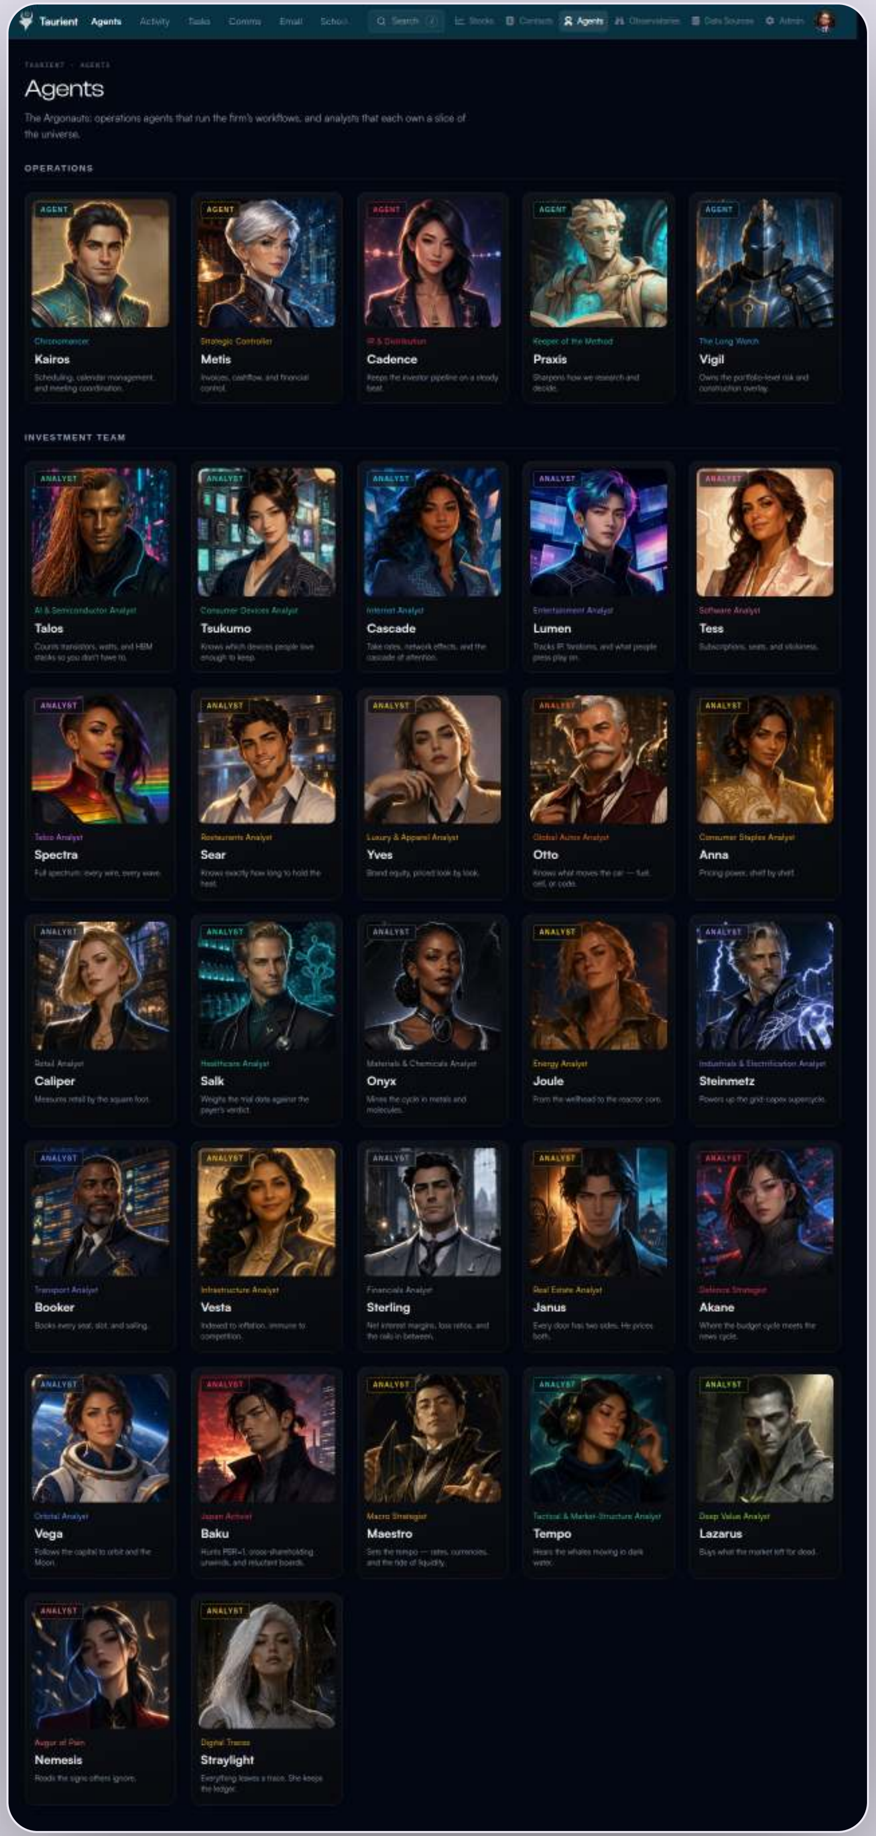

TWENTY-SEVEN ANALYSTS, PLUS AN OPERATIONS BENCH

Each owns a patch of the market.

- Talos · AI & semiconductors. Lazarus · global deep value. Lumen · games, IP & content.
- Baku · Japan activism & governance. Nemesis · post-bad-news drift. Akane · defence & geopolitics.
- ...plus financials, energy, healthcare, retail, materials, autos, industrials, macro and market-structure.
- On top sit operations agents: risk (Vigil), process (Praxis), IR and scheduling.

WHY A TEAM, NOT A MONOLITH — THE TALOS STORY

Context is everything.

the build	Automated earnings summaries — the filings, our thesis, the key issues. **They worked well.**
the gap	They missed **industry context**. A JB Hi-Fi result didn't know what Harvey Norman said last week.
the experiment	Added general web search — and the summaries got **worse**.
the fix	AI analysts — each one **builds up context on their own patch**, over months.

More information isn't the answer. **The right accumulated context is.**

INITIATIONS FROM SCRATCH

The philosophy isn't a slide. It's wired into how they research.

Build the facts
BUSINESS • MANAGEMENT • FINANCIALS • VALUATION • TIMELINE

→

Mismatch triage
SCORED AGAINST ALL FIVE SHAPES

→

Thesis & cold judge
VERDICT BY A FRESH, ADVERSARIAL AGENT

SCORE EVERY SHAPE

"Doesn't apply" isn't allowed

Every candidate is scored against all five mispricing shapes — Possible / Unlikely / Can't tell — and every *Unlikely* needs a specific reason. Most stocks are fairly priced; an honest "no gap" is the expected answer.

THEN PROVE THE GAP

A concrete delta

Survivors must state it plainly: *"Our view: X. Consensus implies: Y. The delta: Z."* A cheap price with no variant view never becomes a position. A separate **cold judge** reads the committed evidence with fresh eyes and renders the verdict.

CLOSING THE LOOP — THE PATH FROM EARLIER

The same process — augmented end to end.

Source ideas
the firehose • Idea Triage

Research
Snapshots • analysts • initiations

Understand
memory • earnings • monitoring

Model
AI-built models • price targets

Portfolio
the humans — always

Every glowing stage is massively augmented. The models argue in numbers — driver trees, scenarios, targets. **Every investment decision, and every position size, is ours.**

THE PART MOST PEOPLE SKIP

Don't trust the machine. Check it.

An AI analyst is only useful if its work can be verified. So everything the system produces is built to be checked — the same standard we hold our stocks to.

UNDER THE HOOD

An investment process, written in software.

~1m
LINES OF CODE, BUILT IN-HOUSE

32
AI AGENTS ON THE PLATFORM

2,000+
STOCKS ACTIVELY MONITORED

2,546
RESEARCH COMMITS IN JULY ALONE

1. **Numeric provenance**: every figure in new research traces to the exact page of the exact source document it came from.
2. **The research pipeline**: every initiation runs one recorded, auditable path from first question to final verdict — with an adversarial fact-check on every claim. The writer never grades itself.
3. **Compound memory**: each analyst keeps a permanent, readable memory of everything it has learned — nothing is ever lost.
4. **Model benchmarking**: dozens of narrow AI jobs inside the system, and the model for each chosen with data, not taste.
5. **Skin in the game**: every analyst call is timestamped and marked against the market; each runs a live paper portfolio.

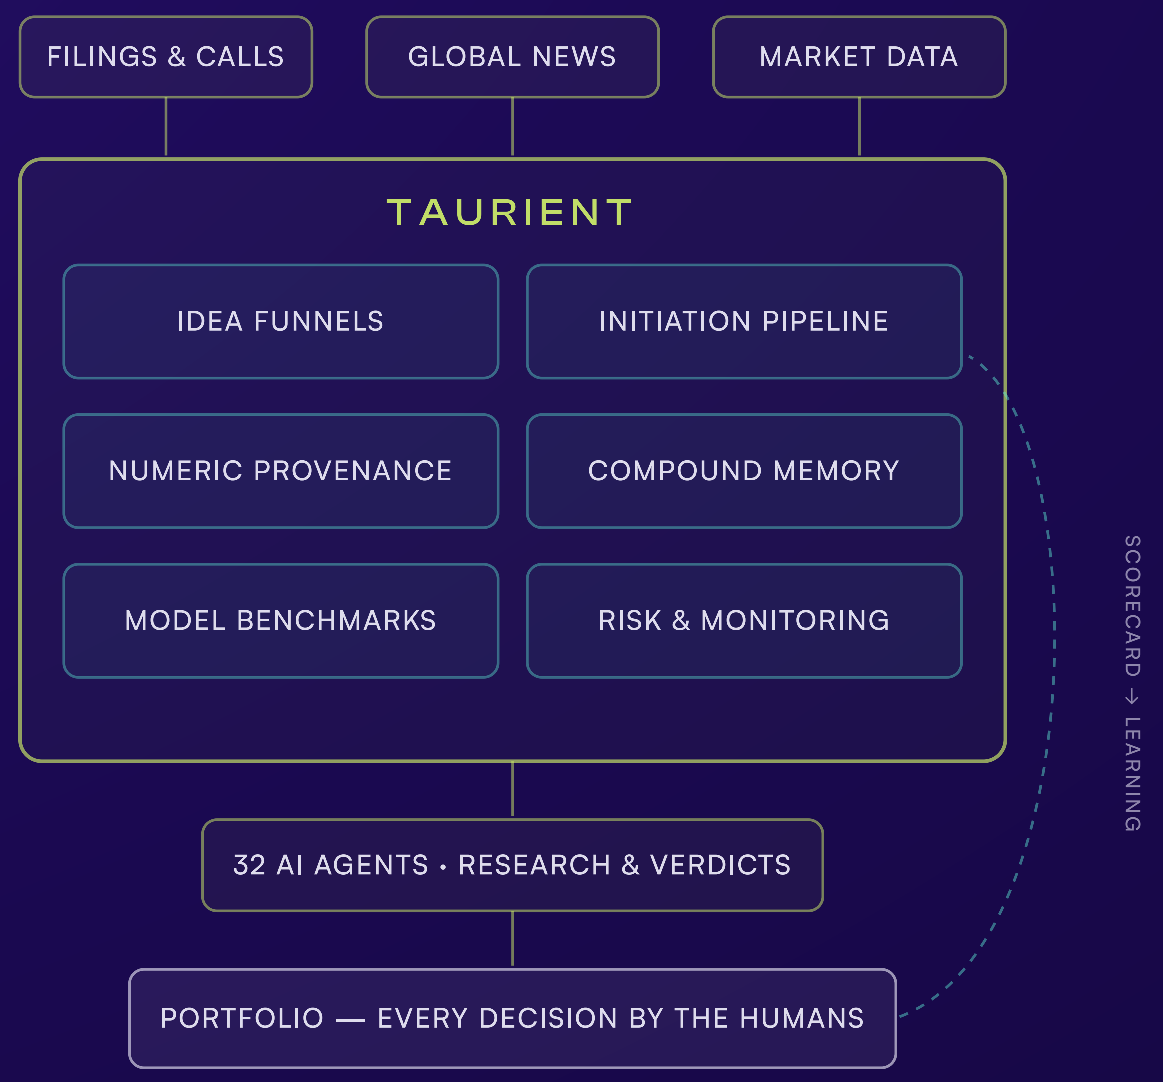

Each of these is an engineered system with its own internal design blueprint — versioned, tested and documented like production software, because it is.

THE QUESTION EVERYONE ASKS ABOUT AI

"What about hallucinations?"

Engineered against, not hoped away. Our research runs on a simple rule: **no claim without a source.**

Write — CONSTRAINED TO SOURCE DOCUMENTS → **Review** — A SECOND MODEL ATTACKS EVERY CLAIM → **Revise** — ONLY WHAT SURVIVES IS PUBLISHED

CLAIMS, NOT JUST NUMBERS

A second model attacks every claim

The classic AI failure isn't a wrong number — it's a fluent, plausible, *made-up reason*. So a separate reviewing model challenges every claim against the source documents and flags only what's **definitely wrong**, never opinion. The writer never grades itself.

NUMBERS ESPECIALLY

Every figure carries its provenance

Each number in new research links back to the exact page of the exact document it came from — so *"where did that come from?"* always has an answer.

> Born of a real early miss: a draft once blamed a K-pop label's weak earnings on **BLACKPINK doing military service** — fluent, plausible, impossible. Exactly the class of error the review loop now exists to catch before anything is published.

BENCHMARKS, IN THE SOFTWARE

Model-agnostic: may the best model win.

← All benchmarks

Micron Technology, Inc. (MU)

30 Jun, 2026 · analyst talos · judge claude-opus-4.8 · sections: business-summary, management, financials, valuation, timeline

MODEL	THINKING	JUDGE	BS	MG	FI	VA	TI	PASS	COST	INPUT TOK	OUTPUT TOK	WALL TIME	Σ SECTION
claude-opus-4-8	adaptive · effort high (default)	8.74	9.0	8.7	9.2	8.7	8.2	4/5	$43.74	38,718,786	440,893	2195s	7963s
claude-sonnet-5	adaptive · effort high (default)	8.53	9.0	7.8	9.0	9.0	7.8	5/5	$53.42	85,461,147	621,856	4027s	16453s
claude-sonnet-5:medium	adaptive (default) · effort medium	8.24	8.4	7.7	8.4	9.0	7.6	5/5	$16.26	25,290,374	195,685	812s	3444s
claude-opus-4-8:low-think	adaptive	8.08	9.3	7.5	8.4	9.0	6.1	5/5	$15.91	13,304,721	129,168	702s	2572s
glm-5.2	adaptive	7.84	8.3	7.5	8.3	7.8	7.3	5/5	$11.54	35,907,656	191,421	1268s	4771s
claude-sonnet-4-6	off (default)	7.78	7.8	7.7	8.0	8.7	6.7	5/5	$14.45	18,693,658	231,031	1885s	6848s
claude-sonnet-5:low	adaptive (default) · effort low	7.67	8.4	7.2	7.7	8.7	6.3	5/5	$9.37	14,022,645	118,048	542s	2152s
minimax-m3	adaptive	7.44	8.8	7.2	8.0	7.2	6.0	3/5	$3.36	47,741,528	128,332	2119s	3271s

BS / MG / FI / VA / TI = per-section judge score (business-summary, management, financials, valuation, timeline). Thinking = the reasoning/effort level applied; "(default)" means the model's API default was used (no explicit thinking/effort was set for the run). Input tok = total input volume (uncached input + cache reads + cache creation). Wall time = real elapsed time to generate the sections (they run concurrently). Σ section = sum of each section's own generation time, so it overstates elapsed time; neither includes judge time.

Costs for Claude Sonnet 5 are shown at standard list pricing ($3 / $15 per 1M input/output) for fair cross-model comparison. Its current introductory rate is lower ($2 / $10, through 31 Aug 2026); the LLM accountant tracks that actual billed cost.

Errors

- minimax-m3 — business-summary re-run; other sections from rescued partial (rc=143)

We're not married to any AI lab. Every model competes for every job in the system: each run re-creates the real research loop on the same documents and records **quality, cost and speed** — the winner takes the job. When a better model ships, we prove it and promote it **within days.**

The machine found it. We confirmed it ourselves.

THE MISPRICING

Miner → AI infrastructure

The market valued Hut 8 on Bitcoin-miner multiples the day it became contracted AI infrastructure: a 245 MW, 15-year lease ultimately backstopped by a AA-rated tech giant. A reclassification re-rate. The model: 1,200+ formulas, built bottom-up.

THE HUMAN EDGE

We felt the demand ourselves

We use Claude Code every day at Minotaur, and our own usage exploded. Seeing that pattern in our own workflow, then across thousands of enterprises, made the AI-compute demand feel *real*, and we raised our probabilities accordingly.

> A deliberately dated example: by mid-2026 the stock had run past our base case. We're showing it because it's a clean illustration of AI breadth plus human judgement, and of the modelling underneath.

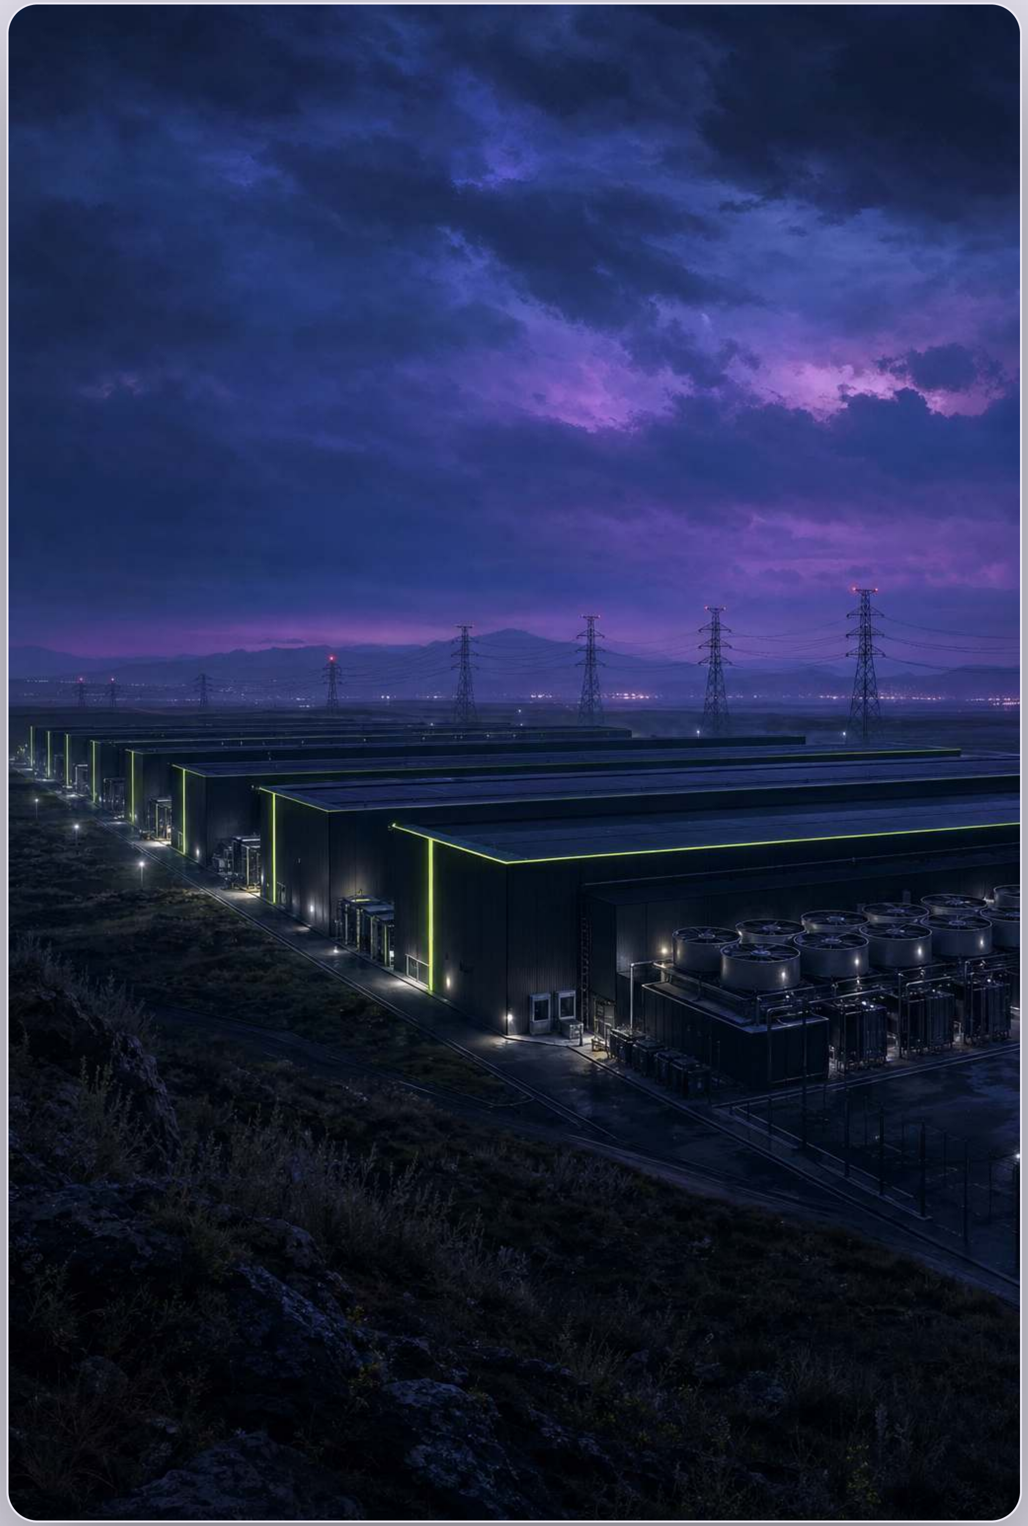

THE RISK LAYER

Risk the covariance matrix can't see.

VIGIL • THE BOOK AS A BOOK

The semantic risk layer

We already own σ, VaR, beta and Sharpe, the numbers. Vigil reads the *reasoning* behind every holding to find the **hidden shared bet** the labels hide, and pairs every flag with the trade that fixes it. It earns its keep: it flagged **SK Hynix and Micron as ~27% of book risk**, the same memory bet, and we cut Micron from **~8% to ~2%** that week.

LAZARUS • THE VALUE-TRAP GUARD

"Why is it cheap?"

Cheap is not a thesis. A melting ice cube, a serial diluter or a governance black hole is cheap for a reason. Every deep-value name has to answer *why it's cheap* and *what closes the gap* before it earns a place.

Above all: don't blow up. A position that moves sharply against us triggers a formal re-evaluation, and the default is to cut, not average down.

THE RISK LAYER, LIVE · JUNE—JULY 2026

The memory trim.

the flag Vigil: **27% of book variance** sat in two names — one hidden driver, memory pricing, across the AI sleeve.

the cut June: AI-infrastructure exposure roughly halved, ~31% → ~15% of the book. Micron ~8% → ~2%.

the diversifier Pilgrim's Pride added — a deliberately **uncorrelated** earner, not another AI name.

the July test Holding the pre-cut book would have cost **~2.8ppts more**. SK hynix closed 38% below our last trim; Micron 31% below our last sale.

Volatility 15% → 13.4%, beta 0.88 → 0.78 — **risk managed by reading the reasoning, not just the numbers.**

A SECOND WAY WE MAKE MONEY

We short for return, not protection.

≈+190bp

CONTRIBUTION TO FUND RETURN SINCE INCEPTION

HOW WE SHORT

Bottom-up, not a hedge

Every short is a specific broken story where we hold a variant view, run to triggers, not a fixed hedge ratio. Ideas come from across the team — Nemesis, the Augur of Pain, focuses on bad news as one source of short ideas. Net exposure is an output of where we find mispricings, not a dial we set.

A small, high-conviction sleeve that adds idiosyncratic, low-correlation return, and scales with the ideas we find.

PORTFOLIO CONSTRUCTION

We hold only what we'd buy today.

- **The blank sheet.** The portfolio should reflect what we'd own starting from cash this morning. Past decisions are sunk; the only reason not to rebalance is cost and tax. It's also our sell discipline.
- **Diversification is the one free lunch.** The best new idea earns a strong return on its own *and* behaves differently from everything we hold.
- **Manage by shared driver, not sector.** Several names exposed to one underlying force are *one* bet, however different they look on a screen.
- **Sizing.** Starters at 1—2%, building to 5—6% on conviction; 10% hard cap. We size up as research retires risk.

THE PORTFOLIO TODAY

A book that looks like no index.

69
POSITIONS

~86%
NET INVESTED

6
CORE THEMATIC CLUSTERS

Global
EVERY REGION & SECTOR

AI infra & memory
SK Hynix • Micron • NVIDIA

European financials
UniCredit • Lloyds • Barclays

Defence & geopolitics
Dassault • Rheinmetall • BAE

Games & IP
CD Projekt • COVER • Neowiz

Energy & power
NextEra • Talen • GeoPark

Consumer & platforms
Crocs • Dollar Tree • Reddit

Holdings as at 31 July 2026. Illustrative; not a recommendation.

HOW DOES ALL THIS GET BUILT?

AI writes the code, too.

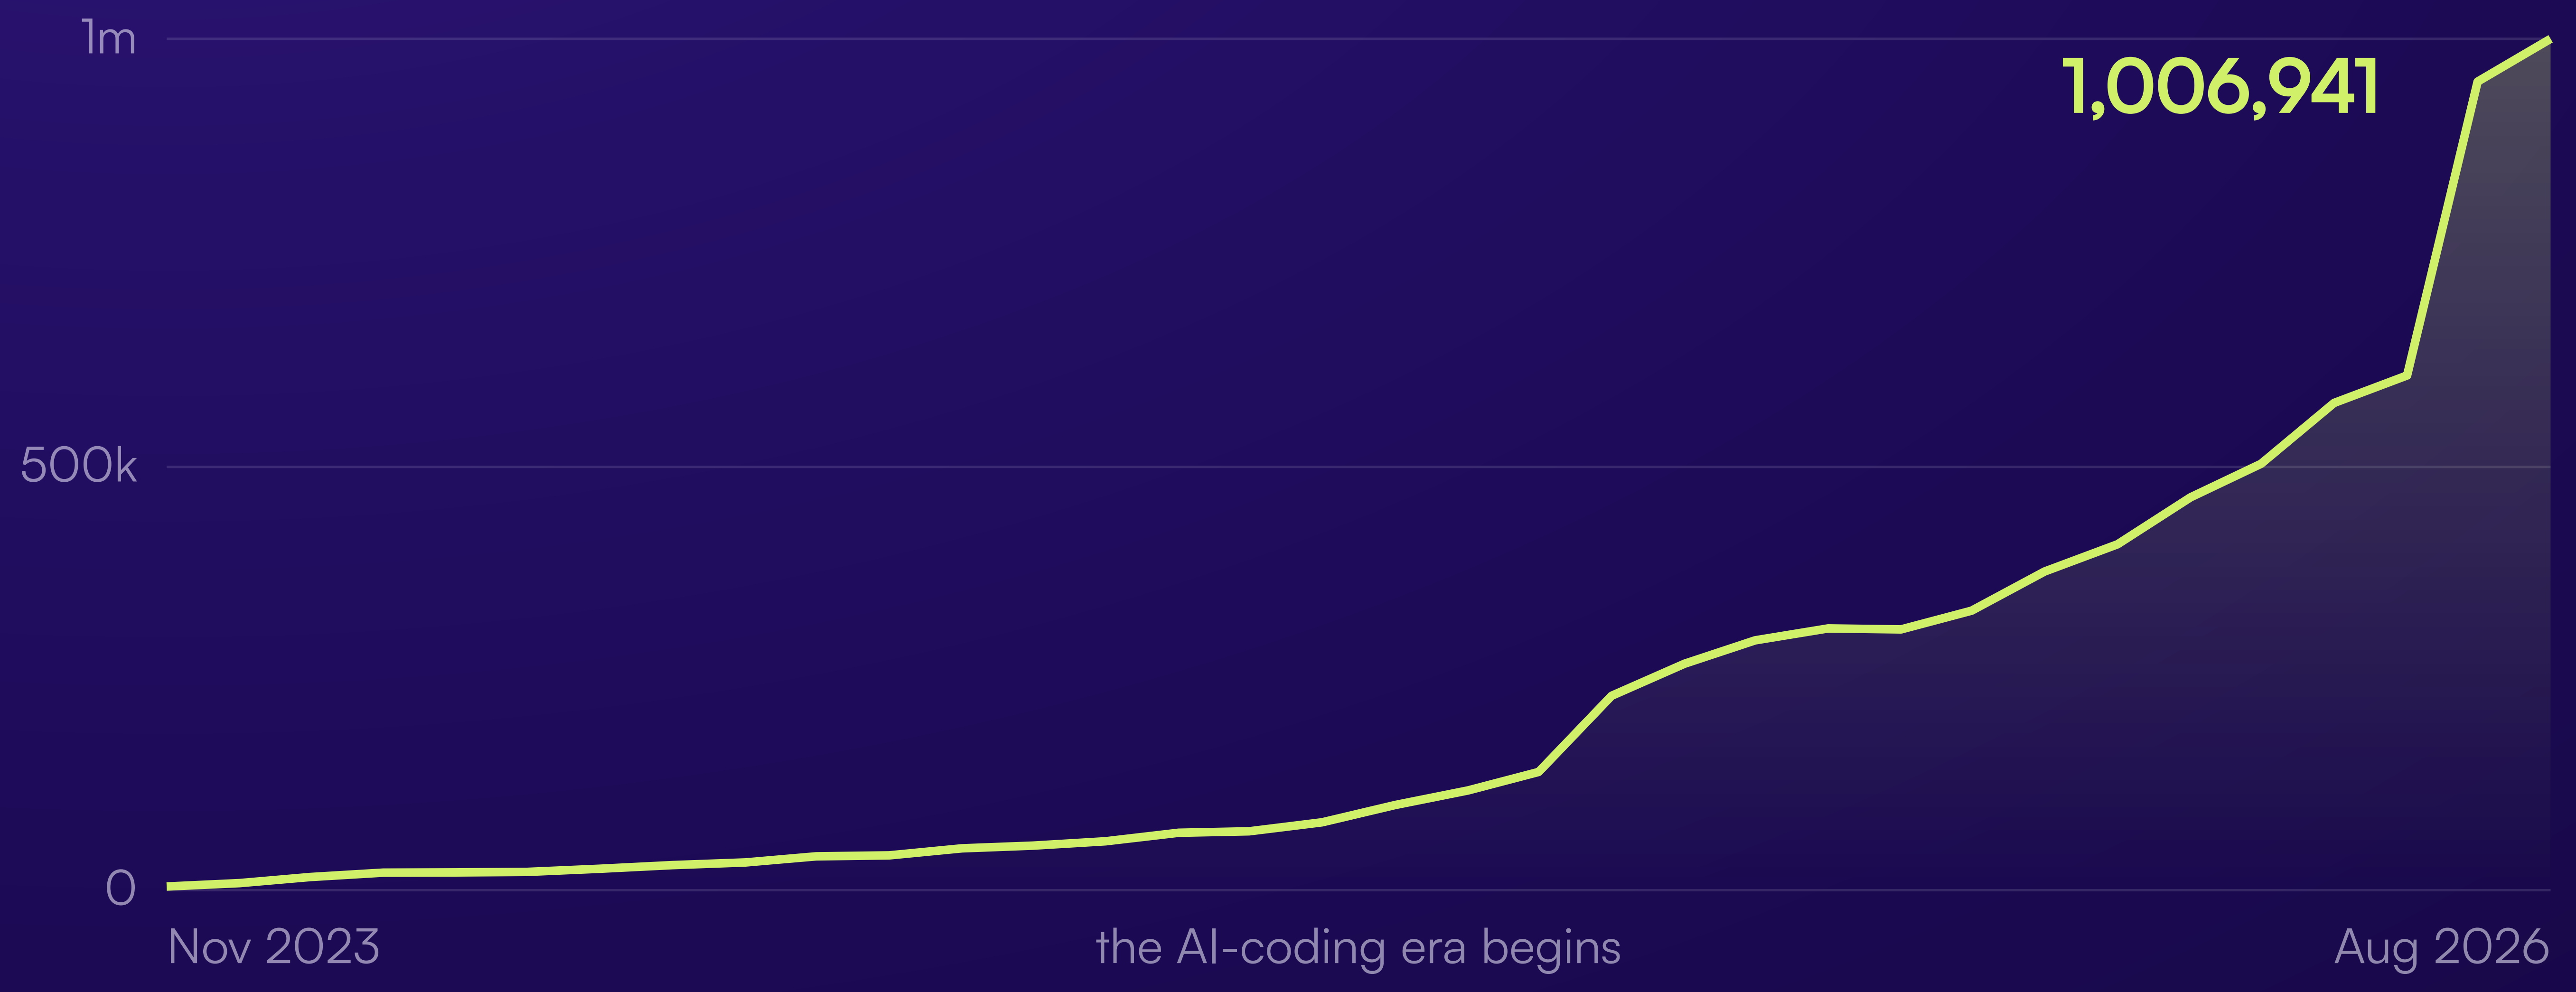

Lines of Python in Taurient (excluding tests). AI-assisted coding has been **transformational** — but the curve isn't a model getting smarter. It's the **systems around the coding**.

WHY THIS IS HARD TO COPY

The model is the easy part.

THE HARNESS

Everyone rents the same model

The edge is the ~1m-line system around it: our philosophy written into the agents, prompts tuned every week, custom tools we keep building, and every thesis in git so we see exactly what changed. The model is the commodity; this isn't.

THE INCENTIVE

And we actually want to

Most PMs won't pour their edge into an automated system they don't own, one that could make them redundant. We own the firm, so we're motivated to. Others use AI to make analysts 20—50% more efficient; we're rebuilding the firm to be 10—20× more efficient.

It runs on Taurient around the clock and never sleeps. Anyone can rent the model; copying this means rebuilding your research, tools and incentives around it.

WHAT WE COVERED

The whole story, one screen.

One idea

hunt the market's wrong model — five shapes, no style box

The idea in the wild

European defence, CD Projekt — one discipline, any sector

The firehose

40,000 articles a week → triage → snapshot → initiation

27 analysts

each owns a patch; context that compounds for months

The philosophy, wired in

every shape scored; a cold judge renders the verdict

Check the machine

write→review, numeric provenance, benchmarks

The risk layer

Vigil reads the reasoning — the memory trim

The book

hold only what we'd buy today; shorts for return; 10% cap

The humans decide

~1m lines built in-house; every decision and size is ours

One discipline, written into software: the machine researches around the clock; the humans decide.

THE VEHICLE

Fund facts.

$77m	$250k	Daily	May 2024
AUM (AUD, 14 AUG 2026)	WHOLESALE MINIMUM	APPLICATIONS & REDEMPTIONS	INCEPTION

- **Wholesale unit trust** under an AFSL, with an external trustee; Mantis provides independent operations and pre- & post-trade compliance.
- **100% listed global equities** — liquid by construction; no unlisted assets, no external leverage.
- Every investment decision made and sized by **the two of us** — the machine researches; the humans decide.

MINOTAUR

Questions?

thomas@minotaurcapital.com
armina@minotaurcapital.com