



MINOTAUR GLOBAL OPPORTUNITIES FUND

Quarterly Report

June 2026



June 2026

The June quarter recovered everything the March quarter gave up and more: the fund returned +16.3%, 2.7ppts ahead of the MSCI ACWI (Net, AUD, after fees). Over twelve months the fund is 3.0% ahead (+20.0% vs +17.0%), and since inception 4.2% a year ahead (+21.7% vs +17.5% p.a.). With the 30 June year-end the fund also paid its first distribution, a small milestone for a young fund. It was a narrow quarter: on our estimates, SK hynix and Micron contributed roughly 15–18ppts between them, and Hut 8, exited in June once it passed our bull-case valuation, added around 4ppts more (local currency estimates before fees and FX movements). That narrowness is exactly what the June de-concentration addressed.


July has been a harder month, with a sharp fall in the memory stocks. The June de-concentration softened the blow rather than avoided it: we kept our highest-conviction memory position at full size, and the fund is currently tracking 1–2% behind the market for the month, with the full numbers to come in the July monthly. Most of the next section is about why we think the fall is an opportunity rather than a verdict.

Memory: what we think, now that it's falling

During July, to the 29 July close, this is what the AI-compute complex did:

A repricing of memory, not of AI

Share-price moves during July 2026, to the 29 July close



Twelve-month moves: SK hynix +432%, Micron +544%, Samsung +187%, AMD +139%, TSMC +91%, Broadcom +22%, Nvidia +6%.

The memory names have fallen by a third to a half during July while Nvidia and Broadcom are down only a few per cent, and the weakness in AMD and TSMC looks like spillover from memory rather than bad news of their own. The selling is concentrated in memory, the corner of the complex that re-rated hardest over the past year, rather than in AI broadly.

What makes July unusual is that the repricing landed on record results. SK hynix's June quarter, reported 29 July, was an all-time high, surpassing the record set the quarter before: revenue of ₩79.3 trillion, up 257% on a year ago, and operating profit of ₩60.5 trillion at a 76% operating margin. Consensus wanted more, and the stock fell 9.6% on the day. Samsung's quarter, confirmed on 30 July in line with its earlier guidance, was also a record, with operating profit up roughly nineteen-fold on a year ago. It beat expectations, yet the shares had already fallen on the guidance day earlier in the month. Record numbers from both, and SK hynix has still lost more than half its value from its 25 June peak, with Samsung about 45% below its own. Expectations had simply run a long way ahead of even record results.


Our memory thesis, for readers newer to it (the fuller version is in the [May monthly](#)): AI systems need memory the way an engine needs a fuel line, demand is outrunning supply, meaningful new capacity doesn't arrive before 2027, and, least appreciated by the market, buyer behaviour has changed structurally, with customers signing multi-year supply agreements rather than buying on spot. In short, we think this cycle runs longer and stronger than a typical memory cycle, and the stocks aren't priced for it: at the 29 July close SK hynix trades at roughly 3x 2027 consensus earnings on its Seoul listing. At a multiple like that you don't need the cycle to last forever, just a little longer than the market's assumption of a short, sharp peak. Capacity constraints alone argue for that.

So why did a record quarter disappoint? Partly, our own thesis working as designed. Those multi-year agreements put a floor under the trough: at Micron the floor is contractual and documented, roughly \$100 billion of contracted revenue at minimum prices under take-or-pay agreements, with some \$22 billion of customer deposits and commitments (around \$18 billion of it cash) behind it. SK hynix's disclosed version, five-year agreements with deposits and purchase commitments but no published floor pricing, is a thinner and weaker form of the same structure. But a floor is also a ceiling: contracted pricing doesn't capture the full spike when market prices run, and analysts cited exactly this as a contributor to SK hynix's revenue miss. There's a second, under-appreciated mechanism: conventional memory currently carries higher margins than HBM (Micron CEO Sanjay Mehrotra and Samsung have both said so on the record) because HBM is priced annually in advance while conventional DRAM reprices quarterly. Samsung expects that gap to narrow materially next year; we think it can reverse when HBM4 repricing lands in 2027. Either way, in a quarter where HBM4 shipped slower than the market hoped, the mix hurt.

The bigger worry sits with the customers, not the suppliers. Consensus estimates for five of the biggest AI spenders (Microsoft, Alphabet, Amazon, Meta and Oracle) have them on track to spend about \$1.58 of additional capex for every additional dollar of operating cash flow between 2025 and 2027, and their combined free cash flow goes from +\$174 billion last year to roughly zero this year on those same estimates (FactSet consensus at 30 July; total capex, since AI-only spend isn't separately disclosed).

Spending the cash flow

Combined capex and free cash flow of Microsoft, Alphabet, Amazon, Meta and Oracle (US\$bn)



Source: FactSet consensus estimates at 30 July 2026, Minotaur Capital analysis. Total capex; fiscal years calendarised as reported.

Companies need to show a return on that spend, and the results in the last week of July were a clean test of which ones can. Microsoft kept its investment plans unchanged, grew Azure 43% in the June quarter as its annual revenue passed US\$100 billion, with contracted future business up 84% (25% excluding OpenAI), and rose about 8% in late trading. Meta grew its revenue 28% but spent just as hard, its free cash flow fell 91%, and without a separately visible AI revenue stream against the spend, it fell. Same night, same AI bet, opposite verdicts. The market is asking to see the return, and we think that's rational.

Part of our research process is writing down, ahead of time, the things that would break a thesis. For memory those are a hyperscaler capex cut, an inventory build at a supplier, or demand normalising in the AI and server channel. None has appeared. [SK hynix](#) raised its 2026 capex by half or more, guided DRAM shipments up 10% quarter-on-quarter into September, and printed an 83% gross margin against the sub-78% level that would concern us. Samsung went further on this week's results call: unmet demand is rolling into next year, it expects the memory shortage in 2027 to be more severe than this year's with tightness persisting into 2028, and since a new fab takes more than three and a half years to reach

production, it sees a significant increase in industry supply before 2028 as unlikely. It also said frontier AI labs, unable to secure cloud capacity, have begun sharing demand forecasts and requesting long-term agreements with Samsung directly, driven by agentic AI token consumption spilling into general-purpose servers, which is the same demand mechanism our thesis leans on.

As AI evolves into agentic forms that perform complex tasks on behalf of users and expands across various services, the underlying demand base for memory is broadening. Consequently, a structural shift is occurring where demand for both AI memory and conventional memory is expanding in tandem.

— SK hynix, June-quarter results release

Two newer narratives deserve a line each, because they're part of what the market is trading. First, the efficiency scare: China's Kimi K3 model launched in mid-July compressing the memory-hungry part of inference, and some read that as AI needing less memory. K3's own launch argues otherwise: its weights alone run to roughly 1.4 terabytes, a rack-scale serving footprint, and demand for the cheaper model was so strong that Moonshot suspended new subscriptions within 48 hours. Cheaper intelligence gets used more; this same fear has been sold twice before in eighteen months, and memory prices rose through both. Second, the Chinese competition scare: CXMT, China's DRAM champion, priced its Shanghai listing at over 300x last year's earnings and then rose 466% on debut, on a sliver of free float. It's a real long-term competitor, but it remains generations behind on cost per bit on our work, and becomes a supply question from 2028, not a reason record earnings deserve a lower multiple today.

What does explain the violence of the move is positioning. Hedge funds added record leverage to global tech through the first five months of the year and have unwound a large share of it since late May. A Goldman Sachs note this week describes one of the largest waves of tech long-selling in a decade, concentrated in Asia, and characterises the move as a deleveraging event rather than a macro one. In Korea, retail margin debt had climbed 140% in eighteen months to a record ₩38.6 trillion by late June, and it has unwound 15% since as the fall forces deleveraging. Selling like that has nothing to do with fair value. It's uncomfortable to sit through, and it's also how a stock in an intact earnings cycle ends up at 3x 2027 earnings.

Which brings us to what we did, and what we're doing. We trimmed SK hynix in late May and cut Micron and the broader AI-infrastructure cluster hard in June (the [June monthly](#) has the full story) for concentration reasons, not because the thesis changed. Holding the pre-cut book through July would have cost roughly 2.8ppts more versus doing nothing (again a local

currency, pre-fee estimate). At the 29 July close SK hynix trades 38% below the price at which we last trimmed it, and Micron 31% below our last June sale. We have begun adding back, buying Micron and Nvidia in the second half of July. Those purchases are underwater as we write; averaging into a deleveraging rarely feels good at the time. But this is precisely what the June de-concentration was for: a diversified book, with the earnings evidence intact and much of the selling driven by deleveraging rather than a change in the facts, leaves us room to buy one of the better risk-rewards we've seen in this cycle.

The risk layer that reads

The de-concentration was Vigil's work, and the June monthly told that story. What we haven't explained is what he adds over the systems we already had.

Our quantitative risk engine (volatility, value-at-risk, beta, per-position risk contributions) does what covariance does: it learns that two holdings share a driver by watching them move together. Vigil is the other half. He reads the *reasoning* behind every position against the live book, looking for holdings that depend on the same thing while sitting in different sectors. The numbers can only tell you two holdings share a driver after they've moved together. Reading the reasoning behind each position can tell you before they do.

The June demonstration: a conventional sector view splits our AI-linked holdings across four different buckets: memory, semis, data-centre power, and software names that shared the same capex-cycle dependency. No single label showed the full bet. Reading what each holding actually depended on did, and that reading is what sat behind the decision the [June monthly](#) describes: halving AI-infrastructure exposure from roughly 31% of the portfolio to around 15%. The same pass questioned whether our precious-metals sleeve was the diversifier it looked like; we largely sold it in the June de-gross.

Vigil operates under hard rules written into his standing instructions: alerts must carry a proposed action; silence is the default; the quantitative numbers are never overridden, only explained and challenged.

Vigil's dossier: his portrait and operating rules, as they appear in Taurient · tap to enlarge

By end-June the de-grossing had done what it was meant to: volatility fell from roughly 15% to 13.4% annualised (from above the index to below it) and beta from about 0.88 to 0.78. Vigil's note on 27 June then put a number on the trade-off: trimming broadly while leaving SK hynix in place had lifted its share of measured risk to roughly a quarter of the total. The conviction

case had been formally re-underwritten the same month, so it was a choice, and July tested it. Every concentration in the book should be a decision rather than an accident; that's what he's for, and it's also why we can add into this fall rather than sell into it.

The quarter the tool became a team

Three months ago we told you Taurient had reached the point where we could research a whole industry at once rather than one company at a time. That was true, but it described a *tool*. Since then the tool has become something closer to a *team*: a roster of named AI analysts, each with its own coverage universe, its own memory, its own reading list and its own email address, arguing with each other and with us.

Talos was the first, introduced in the [April monthly](#) as an experiment. By the [May monthly](#) there were 19. Janus, who covers listed real estate, joined on 12 July and made it 23, plus five operations agents. The roster could grow that fast because we didn't build 23 bots. We built one learning-analyst platform and configured it 23 times. Each analyst gets the same machinery:

- a coverage universe and a set of subscribed sources it reads on its own schedule,
- durable memory, where every change to a view is dated and attributable,
- a monitoring loop that wakes it when something it covers moves.

What differs between analysts is custom code. Baku reads Japanese large-shareholder filings straight from EDINET, Vega reads FCC spectrum filings, Anna reads NOAA weather and USDA crop feeds. Each is purpose-built software behind that one analyst, quick to build because Taurient is the scaffolding it all plugs into. Increasingly the agents propose their own. In May, Vega suggested she should be able to count satellites – actual hardware on orbit, from public tracking data, as a check on what constellation companies claim. An AI coding agent built it, a second AI reviewed it, and Thomas did the final review, catching a design error both AIs had missed, before it went live. She now gets a daily feed of the real deployment cadence of every constellation she covers, and has since proposed an improvement to her own tool – with us as the final gate on all of it.

The agents also consult each other. When Vega needed the defence view on a space company's missile-defence bookings, she didn't get a generic answer. She got Akane's actual position, grounded in her stored research, separating the real program of record from the aspirational one and naming the margin risk in between:

A desk query: Vega asks Akane about Voyager Technologies' Golden Dome bookings · tap to enlarge

The value in that exchange isn't a second opinion. Our analysts largely run on the same underlying models, and a model debating itself proves little. A [University of Maryland study](#) posted in June measured this: across more than a thousand human contributions to long-form public debates, 65.3% of the main arguments people made were unique, while for LLM-written essays the figure was 3.4%. (The domain was op-eds rather than stock research, so read the transfer as our inference.) When Talos and Steinmetz reach the same conclusion, that's one opinion drawn twice. What makes a desk query worth sending is context, not opinion: Akane covers defence full-time, reading the results, news flow and budget debates that Vega never touches, so her answer carries information the asking analyst doesn't hold. The same logic runs through the research process: agreement between two of our agents isn't treated as confirmation, every initiation must produce an adversarial brief (the strongest available case against its own conclusion), and a "cold judge", a model from a different family shown the evidence but not the reasoning, grades the conclusion blind.

The output is now the bulk of the firm's research volume: in July so far our analysts have produced earnings summaries on 103 results across 101 covered companies, more than double June. A new modelling system also went live during the month, the most substantial upgrade to our modelling capability since the firm started, with 15 house models already rebuilt under it. More on that in a future report. The pace of building behind all this is the fastest it has been, though pace, like every input measure, only counts if it shows up in fund performance, which is the only test that matters.

The May monthly [promised more depth](#) on the operations agents. Watts Humphrey wrote two decades ago that every business is a software business. Most still aren't, but they should be. We thought about that when starting Minotaur and built it as one from day one. Taurient is effectively the firm's operating system: the CRM, the research repository and the portfolio live inside it, and what doesn't (email, the calendar) it reaches programmatically. An operations agent bolted onto tools that can't see each other is a chatbot; one that hooks into the firm's OS can do the job. Three rules bind the ops agents: judgement, not just tasks; nothing crosses the firm boundary without a human; answer from systems of record, never from memory. Kairos runs the calendar. When he proposes meeting times he places provisional holds in the same pass, and before a trip is locked he geocodes the venues, runs the drive times and builds the travel buffers in. Metis, our financial controller, answers contract questions from the executed agreement itself: a date with a consequence comes from the primary document, not the summary record. And when an investor asked this week for our monthly return series and a longer benchmark history for a correlation study, Cadence built the workbook, flagged that the part-month at inception would distort his volatility maths, left out a July estimate that isn't an official number yet, and pointed to the clean source for the benchmark history she doesn't

hold rather than improvising a series. Nothing she drafts reaches an investor without one of us sending it. Praxis owns the investment process itself; he reads the frontier of AI research and folds what's useful into how we work.

A common question on the road this quarter was some version of: "fine, but how do you know any of this is adding value?" It's the right question. The roster is a few months old, so none of the fund's track record can be attributed to it – until June the answer rested on judgement rather than data, and in June we built the measurement. Since 20 June our AI sector analysts each maintain a market-neutral conviction book across their coverage: long and short sleeves each normalised to 100, three levels of conviction on each side, entries append-only so every revision is dated and preserved. Five weeks of data is nothing, so there are no results to show yet. The point isn't to rank the agents and cut the losers – it's to find which analysts need better instructions, better sources and better skills, and fix them.

On cost: our model spend now runs at over A\$10,000 a month and is rising as the roster works harder, still far less than a single junior analyst. That's model spend only; it excludes data subscriptions and our own time building the system, which is the real investment and doesn't show up on an invoice. Cheap isn't the argument anyway, as an adviser in Sydney reminded us: spend is an input, and the scoreboard above is where the output will show.

So what is an analyst now?

A [working paper](#) by an AllianceBernstein strategist argues that asset management is going through a "value-stack inversion": information gathering, analysis and first-pass portfolio construction are commoditising fast, so the scarce resource stops being any individual forecast and becomes the architecture by which a firm generates a belief, challenges it, sizes it, and learns from being wrong. If that's right, the metrics that matter are strange ones: not "how many companies do you cover" but how fast a firm revises a belief once disconfirming evidence lands.

The same paper names the failure modes of what we're building. The most dangerous is "synthetic coherence", the shared-model problem from earlier. A roster of AI analysts that all quietly agree with each other, faster and faster, is a worse research process than two people arguing. It's why the adversarial brief and cold judge processes exist.

Here's the question we find more useful than any benchmark. Would Salk, our healthcare analyst, beat an experienced biotech specialist today? No. Not because the models lack intelligence, but because they don't yet know what it takes to outperform the market in biotech: which trial readouts matter, which management claims to discount, where the bodies are buried. Our job is teaching them, and unlike training a human analyst, the lesson sticks, across every stock.

That leaves a question a few of you have asked directly: if you can run a whole team of AI analysts for the cost of one junior analyst, why would you ever hire a person? For two years we've said Minotaur is two portfolio managers with no analysts by design, and that thinking has changed: the roster now surfaces more well-evidenced ideas than two people can act on, so the constraint has moved from research production to decision bandwidth. The role we'd hire for looks nothing like the analyst role of 2023. It's someone who works on the system: orchestrating a team of AI analysts, validating their output, and teaching their domain expertise into desks that never forget it. When we wrote in December 2024 that AI was our superpower and people were still at the heart of Minotaur, we meant it.

Closing thoughts

The June quarter was a good one, and July has been a hard month for our largest theme. The earnings evidence hasn't turned; positioning has; and the diversified book we built in June is what lets us treat the difference as an opportunity. The triggers that would break the thesis are written down, and the scoreboard on our analysts is running.

As Solon said, "I grow old ever learning many things."

That's roughly the job description now.

READ MORE

Explore the interactive report

ONLINE VERSION

Charts, images, and supporting links

The web edition of this report has all the supporting charts and any videos or embedded media referenced in the commentary.

<https://www.minotaurcapital.com/reports/quarterly/2026-06>

For investor enquiries, contact invest@minotaurcapital.com.

IMPORTANT INFORMATION

Minotaur Capital Management Pty Ltd (ABN 17 672 819 975) is a corporate authorised representative (CAR 1308265) of Minotaur Licensing Pty Ltd (ABN 86 674 743 198) (AFSL 557080). The Minotaur Global Opportunities Fund is issued by K2 Asset Management Ltd (ABN 95 085 445 094, AFSL 244393), a wholly owned subsidiary of K2 Asset Management Holdings Ltd (ABN 59 124 636 782).

The information in this document (the Information) has been prepared by Minotaur.

This information is for general information only and is not an offer for the purchase or sale of any financial product or services. The Information has been prepared for investors who qualify as wholesale clients under section 761G of the Corporations Act 2001 (Cth) (Corporations Act) or to any other person who is not required to be given a regulated disclosure document under the Corporations Act. The Information is not intended to provide you with financial or tax advice and does not take into account your objectives, financial situation or needs. Although we believe that the Information is correct, no warranty of accuracy, reliability or completeness is given, except for liability under statute which cannot be excluded. Please note that past performance may not be indicative of future performance and that no guarantee of performance, the return of capital or a particular rate of return is given by Minotaur, K2 Asset Management or any other person. To the maximum extent possible, Minotaur, K2 Asset Management or any other person do not accept any liability for any statement in this Information.