



MINOTAUR GLOBAL OPPORTUNITIES FUND

Quarterly Report

December 2025



December 2025



AI Has Levelled Up... Again

We know we talk about AI a lot. But the pace of acceleration in the final weeks of 2025 surprised even us – enough that this quarter we’re dedicating most of our commentary to what’s changed, what we’ve built, and what it means for software stocks.

Two shifts have converged in the last few months to create what feels like a step-change in AI capability. Neither is purely about the models getting smarter (though they have). Both are about how we use them.



Shift One: Skills Expand What AI Can Touch

The first shift is skills. A large language model can only take so much text (known as the context window), and the more you give it, the worse it tends to perform. If you want it to know how to do a hundred different things, you could give it all those instructions upfront, but that degrades performance. Skills solve this by giving the model a short description of what it could learn upfront, and if it needs that capability, it asks for the full instructions. It keeps the model focused and working better.

When Anthropic launched their Skills system in [October](#), they provided a standardised format for packaging and sharing these instructions. The [Skills Marketplace](#), a collection of skills aggregated from open sources (not vetted, so read any before you use them), now has over 77,000 skills. Your agent doesn’t know how to work with Excel? Give it a skill and now it does.


We’ve built our own library of skills that let the agent access external services and take actions on our behalf. A lot more tasks are now automatable, and because skills are shareable, the ecosystem compounds.



Shift Two: From Chatting to Loops

The second shift matters more. We’re moving from *chatting with AI* to *running AI in loops*.

Geoffrey Huntley, previously Canva's AI Developer Productivity Tech Lead, popularised a deliberately unsophisticated approach under the name "[Ralph Wiggum](#)": define what "done" looks like in a specification file, run the agent in a loop, give each iteration a fresh context window but let it see what was built before. As Huntley put it when explaining the name: "It's kind of dumb, kind of lovable, and it never gives up." The contrast with complex multi-agent orchestration systems is striking: the dumb approach, through sheer iteration, often outperforms the clever one.

A terminal window titled "ralph - building-crm" with standard macOS window controls (red, yellow, green buttons). The terminal shows the command "\$ ralph --spec 'Build a modern CRM from scratch'" being executed. It then displays "Loading PROMPT.md specification..." and "Target: Full CRM with contacts, companies, deals, activities, reporting". A separator line is followed by "Iteration 1". The goal is "Set up Next.js project with TypeScript", and the command "npx create-next-app@14 crm --typescript --tailwind --app" is shown. At the bottom, a green dot and the word "Running" are on the left, and "Iteration 1" is on the right.

```
ralph - building-crm

$ ralph --spec "Build a modern CRM from scratch"
Loading PROMPT.md specification...
Target: Full CRM with contacts, companies, deals, activities, reporting
— Iteration 1 —
Goal: Set up Next.js project with TypeScript
npx create-next-app@14 crm --typescript --tailwind --app

• Running Iteration 1
```

The Ralph Wiggum technique: run AI in a loop until the task is complete.

This one is building a CRM (customer relationship management system). Accelerated for dramatic effect.

Huntley first presented it at a San Francisco meetup in June 2025, published the approach in July, and by December it had gone viral. Anthropic released an official Claude Code plugin for it. VentureBeat ran a [feature story](#).



geoff ✓
@GeoffreyHuntley



tired:

overpaid sf engineers doing complex shit

wired:

some australian doing bush engineering



3:14 AM · Jan 12, 2026 · **6,512** Views



10



9



184



21



Geoffrey Huntley on his unconventional approach to AI development.

What made loops suddenly viable? As Huntley [put it](#): “Over the last year the models have become quite good and it’s only now that I’m able to realise this full vision.” Claude Opus 4.5 arrived in late November; GPT-5.2 and GPT-5.2-Codex followed in December. Better harnesses helped too. A harness is the layer between you and the model. It’s what lets the AI read files, run commands, and take actions rather than just chat. Tools like Claude Code, Cursor, and [OpenCode](#) handle context management, permissions, and multi-session workflows, making the models far more capable.

What We Built

We experienced this over the Christmas break. Thomas built a personal AI assistant (and a Minotaur Assistant for the team) with skills for email, calendar, Slack, Xero, Confluence, web search, image generation, and our internal fund reporting. He also built automation scripts that run AI agents in loops against task lists, converting unstructured notes into structured plans and executing them overnight while he slept.

Now we can give the assistant tasks like “read this email from our lawyers and update our notes on ETF setup costs” or “why is Lenovo down?” (which triggers a skill that simultaneously searches news, X, and email before responding). It can read an investor’s due diligence request, pull the relevant fund data, and generate a customised Excel report tailored to the specific request – a task that could previously take hours of manual work. Each new skill expands what’s automatable without polluting unrelated tasks with irrelevant context.

What This Means for Software

This brings us to a debate we’ve been having internally. We suspect many investors are having it too.

A few days after Thomas's holiday coding sprint, our exchange went something like:


Thomas: "I'm not sure about software after my week of coding..."

Arms: "I think Atlassian will be fine."

Thomas: "There's been a step change... my week just made it visceral."

That word *visceral* matters. These improvements don't arrive as neat, linear progress you can tidy up in a spreadsheet. They arrive as sudden “oh wow” moments that change what feels possible.

Software stocks have been hammered. HubSpot is down over 55% from its January 2025 levels. Atlassian has more than halved. Adobe and Salesforce are both down around 30%. Oppenheimer downgraded Adobe in mid-January, stating software has “flipped from an AI beneficiary to an AI victim.”



“I’m in danger.” – Ralph Wiggum

Part of it is sentiment catching up to what developers have been experiencing. If AI can compress the time and cost to build software by an order of magnitude, what happens to the companies that sell it?

There’s an economic shift at the heart of this. A year ago, AI wasn’t good enough to substitute for software development; you had to use developers. Now the output quality has crossed a threshold where it’s genuinely useful, and it’s cheap. If capable LLMs were free and infinitely fast, you’d always prefer brute force iteration over careful human design. We’re moving in that direction. The dumb loop that runs overnight starts to look more attractive than the carefully reviewed pull request.

What does this mean for software stocks? We think about the risks through three lenses:

1. Build vs Buy

Can organisations code their own solutions instead of buying software? For some tools, yes. But there are reasons most won’t: maintenance burden, liability, what happens when things break at 2am. That said, complex enterprise software where you only use 10% of the features? Sometimes a custom tool that does exactly what you need, integrated exactly how you want, is genuinely better. It’s one reason we write a lot of software ourselves.

2. Competitive Intensity

If the cost and time to build software falls materially, what happens to competitive dynamics? New companies can spin up faster. Existing competitors can ship products faster. The cost to

build "good enough" alternatives falls. This doesn't mean incumbents lose, but it lowers the barrier to entry, which changes who can afford to compete.

3. Per-Seat Pricing

Because AI can do more jobs, particularly entry-level tasks, we wouldn't be surprised if teams simultaneously produce more while being smaller. Services that charge per seat face pressure. Transitioning to different pricing models (outcomes-based, consumption-based) isn't trivial.

None of this means "software is dead." The simplistic "AI kills software" narrative is too linear. But the risk environment has changed, and it makes sense that these stocks are re-rating.

Even sceptics are shifting. David Heinemeier Hansson, creator of Ruby on Rails and a vocal sceptic of AI hype, recently admitted that half his resistance was "simply that the models weren't good enough yet." That's no longer the case.



DHH   @dhh · Jan 4



You can't let the slop and cringe deny you the wonder of AI. This is the most exciting thing we've made computers do since we connected them to the internet. If you spent 2025 being pessimistic or skeptical on AI, why not give the start of 2026 a try with optimism and curiosity?

 209

 513

 4.9K

 303K



DHH   @dhh



Just this past summer, I spoke with [@lexfridman](#) about not letting AI write any code directly, but it turns out half the resistance was simply that the models weren't good enough yet! I spent more time rewriting what it wrote than if I'd done it from scratch. That has now flipped.

4:27 AM · Jan 4, 2026 · **342.5K** Views

 41

 182

 1.8K

 371



Even AI sceptics are updating their views.

The Orchestration Layer

One dynamic that interests us is where value accrues in this new stack.

If an AI assistant becomes the layer that sits between you and your services – email, calendar, documents, CRM, research, analytics – then the underlying UI becomes less important. The “winner” is whoever owns the permissions, the workflow graph, the audit trail, and the ability to safely take action across systems.

That has second-order implications. Some software becomes *more* defensible (if it is the control plane). Some becomes *less* defensible (if it's a thin UI over commoditised functionality). Value may shift toward platforms that can govern agentic workflows rather than those that just serve screens. We're already living this. We don't pay for software unless we can access it via an API. If it can't fit into our automated systems, we don't use it.

On January 12, Anthropic launched [Claude Cowork](#), a desktop agent that manages files, connects to services like Notion and Asana, and executes multi-step tasks autonomously. It's essentially a consumer-friendly version of what we built ourselves over Christmas, though less powerful. Microsoft is integrating Claude into Copilot. OpenAI is pushing ChatGPT as the default interface. The race to own the orchestration layer is already underway.

Portfolio Implications

In terms of what this means for our portfolio, we've reduced our software exposure since quarter end, including cutting Atlassian, despite our bullish commentary on the name and sector late last year. We sold because the goalposts are shifting fast on unit economics and defensibility of seat-based workflow software as agentic tooling spreads. The distribution of outcomes has widened materially, and at current prices we don't think the risk-reward is attractive.

We still believe in the AI supercycle. Our positions in Nvidia and the memory makers (discussed in our [December monthly](#)) reflect this. Compute and infrastructure demand remains strong.

We also initiated a position in Hut 8 during the quarter. On December 17, Hut 8 announced a 15-year, 245 MW data centre lease to Anthropic at their River Bend campus in Louisiana, backstopped by Google (AA+ credit), with a total contract value of US\$7 billion. We bought the stock that day. We were able to move quickly both because of AI-assisted analysis and because the deal structure was familiar: in a prior role, Thomas was a large shareholder in NEXTDC and Asia Pacific Data Centre, the trust NEXTDC created to fund their land and buildings when capital was scarce. River Bend has a similar feel – contracted, investment-grade-backed infrastructure with visible cash flows. At mid-

point build cost estimates and before capitalised interest, the project yields ~15% unlevered in year one, rising with 3% annual escalators over the lease term. We thought the market was underpricing the impact of this deal on the company.

Staying Right as Facts Evolve

In a world that compounds change, good investing isn't about being right once. It's about staying right as the facts evolve.

AI moves in dog years. The past twelve months feel like seven years of progress. OpenAI's Deep Research isn't even a year old. Skills launched three months ago. Opus 4.5 arrived two months ago. The Ralph Wiggum pattern went viral this month. Each felt like a step change at the time. They're stacking up. The big unknown for software stocks isn't where we are now. It's where we'll be in another twelve months if this pace continues.

We see this in our own work. Since starting Minotaur, AI has become increasingly central to our research process. With this latest wave, it's now thoroughly entrenched. We're building AI systems that are deliberately "opinionated", in that they encode our philosophy, our decision rules, our view of what matters. The problem with using off-the-shelf ChatGPT to research or summarise is that by the time you're looking at the output, you're trying to extract insight from something that's already been compressed. Information has been lost. The signal you needed might have been discarded as noise. Because we've infused our AI with how we actually think, we can apply that lens at the detail level, before the lossy compression, not after.

The progress we've made incorporating AI into the deeper parts of our research process makes us think it won't be long before we can generate full investment recommendations with AI. At that point, we'll be able to measure whether the human judgement of Thomas and Arms adds or detracts value beyond the systems we build.

Diogenes once said to a would-be philosopher about the importance of lived experience vs. textbook: *"You do not choose painted figs, but real ones... and yet you pass over the true training and would apply yourself to written rules."* In markets right now, there's no shortage of "painted figs", ie. investors opining loudly on AI's impact on software without actually using the tools. Our advantage is simple: we eat the real figs. We build with these systems, run them in loops, and rely on them both in research and operationally. That's why we're not guessing when a step-change matters – we actually feel it, which is why we can change our mind with discipline when the facts move.

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/2025-12>

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.