

INVESTMENT COMMENTARY

Q1 2026 | JUNE 10, 2026



Commentary Highlights

OUR AI WORKFLOW

—
ARCHITECTURE & STRUCTURE

—
WHAT NEXT?

—
TURNING PROCESS INTO INSIGHTS

—
AMAZON / AWS THESIS

—
SAAS RISK TRACKER

Year Zero: How AI Is Reshaping Our Investment Process

In our last commentary, we discussed Claude Code as “a more recent discovery” with “jaw dropping” potential. With the benefit of hindsight, we were too restrained in sharing our enthusiasm for Claude. In many respects, the past few months have felt like “Year Zero” for our research process, reorienting and rebuilding our tools with and around Claude Code. Chatting with LLMs is helpful, but we have learned this year that it only scratches the surface of AI’s real potential. By leaning into these new discoveries, we have not just enhanced our process, we have already generated key investment insights.

The key point is not that AI makes research faster, though it does. The more important point is that it changes the surface area of what we can monitor, test, and revisit. We can now track more companies, more inputs, and more changes without diluting the quality of our attention. For an investment process built around patience, selectivity, and evidence, that is a meaningful change.

*Building long-term portfolios
through deeply researched
investments.*

RGAIA.COM

In this commentary we will walk you through our workflow and then discuss a few specific investments. We have several goals in sharing some of our discoveries here:

- We want you to share in our enthusiasm for these new tools.
- We hope others will share ideas with us on how we can become more productive and generate even more value.
- We want to hold ourselves accountable and track our progress over time. We cannot yet quantify the ROI of these tools, but we expect to demonstrate their value more objectively over time.

At the outset, we should be clear that we do not view the deployment of AI itself as proprietary, though we have built proprietary tools that we will not share or discuss. AI is the ultimate force multiplier for human thought, it is not a replacement. Our process is fundamental to our ethos and does not change. RGA believes deeply in a low turnover, GARP orientation, with an appreciation for quality. In fact, in our version of Claude's core memory (Claude.md), we have memorialized our own investment worldview with a memo describing our workflow from idea generation through portfolio management. Stated another way: Claude, as we use it ourselves, is deeply indoctrinated in our worldview and process. The only real change is in the tools we are using. We have replaced Factset with a combination of AI models and a handful of APIs.

OUR WORKFLOW

The workflow starts with our dashboard that pulls together all of our various projects. At the very top, we see the results of our agentic project manager. If any of our various projects fails to run or launch as expected, we get a large red tile indicating which project we need to troubleshoot and why it failed to run. Given the time we have put into building these projects, those failures are increasingly infrequent, though it is incredibly important to know if what you are looking at is clean, factual information or something is broken.

Next, we have embedded links into each of our projects, sorted by category:

- Interactive Tools
- Expert & Management Calls
- Industry Dashboards
- Consumer Demand Trackers
- Filings and Macro
- Screening and Quantitative
- Live Signals
- Alternative Data
- Company Deep Dives
- Cross-Project Synthesis

Cross-Project Synthesis then leads into the next section: our daily "Cross-Project Memo." Each day, this memo focuses on the critical changes across all of our dashboards. Change here is the key—we isolate and focus on what new material surfaces and where there are notable deltas across our various projects. In the delta lie the insights and the questions that we need to pursue. The cross-project memo is heavy on bullet points and visuals. It flows into/concludes "What to watch" and "suggested follow-up inquiry" and these are geared to our North Stars. In a similar fashion to all of our workflows—it concludes with a hole finding agent which uses

a cross model validation framework we have developed internally to identify gaps and potential data weaknesses and/or misstatements.

Beyond these projects, we have built dozens of skills. These are not skills with AI, but rather ones that we have taught Claude and built into repeatable workflows. We have built dozens of skills ranging from more rudimentary pieces of our own workflows to call prep and synthesis. These skills are not exactly projects per se, but they help feed into them and have been tremendous accelerants in our own process.

ARCHITECTURE AND STRUCTURE

Our preferred setup uses the Antigravity IDE, Google's AI-native integrated development platform. We leverage a number of tools therein, with Claude Code in the command-line interface, Gemini side agents for efficient planning, large context windows, and efficient token use, and Codex for heavy quantitative validation work. We are also increasingly leveraging Claude Code via the desktop app for simple recurring tasks we share across our team.

Ryan's background in CLI has been extremely helpful from the outset. As a consequence, early on, we developed an appreciation for building with thoughtful, scalable architectures and optimizing for token efficiency. These are critical points that ultimately have significant investment ramifications, but should also be at the heart of how projects are executed. Although this commentary focuses on what we have built with Claude, most of these projects do not run on Claude. Most of our projects run on Python scripts, rely on APIs that access structured information, and do not use AI during execution. We have simply leveraged Claude to build durable tools. To the extent they do utilize AI, the LLMs are accessed via API calls with specific parameter settings which we have refined internally. We further cache all the resulting LLM outputs in our own databases—allowing for cost efficient future access. These projects use SQL, JSON and MD files and we have chosen to visualize them in locally stored HTML files.¹ Claude knows our preferred architectures and folder structures, so each project we start immediately builds in exactly the same, organized way.

As mentioned, several of these projects leverage AI along the way, and for that, we tap into the LLM model of our choice. This is an important point that you will hear more of over time. Although we are mainly using Anthropic's Opus model via Claude Code to build, we are carefully selecting the appropriate model for the given task. For more rudimentary operations (think aggregating numbers, more like data entry), we are using the cheapest capable model and for more complex analyses that are semantic in nature, we tend to use Gemini. For numerate and internal reconstruction, we use Sonnet or Opus. These simple rules of thumb are subject to change, but model token efficiency aside, we expect our actual use of AI, as measured by the volume of tokens we burn, to level off or even decline once our phase of heavy building is behind us.

WHAT NEXT?

The beauty of this structure is that our projects are evolving into the RGA Investment Management Operating System. Our thesis is housed in our own words, nested within how our ideas are being tracked. We are building structured datasets across key areas of our work, and while we are already harvesting insights, we expect the output to grow meaningfully over time. This is happening in a variety of ways. Each of our screeners is built with a real-time performance tracker; in other words, we will objectively know which screens generate value and which ones do not.

1. We will soon migrate everything to a secure, virtual host as our primary portal, but that's not exactly necessary today.

To share a few small examples—we are acquiring data points on key inputs for our companies, ranging from points of distribution to pricing to sentiment of reviews and tracking the progression over time. We have developed the logic that sits behind these workflows and translates this data into actionable insights with quantifiable signal value.

Soon, we will undertake a critical step forward, which was referenced above. We will be nesting our projects in a private domain, where our key assets are hosted and accessible online. We may instead forego online hosting and use a physical server that we can access directly. If you are reading this and have strong opinions on the functionality and security of either path, please do let us know.

The most important test for these tools is whether they lead to better questions, analysis and insights. Dashboards and agents are only useful if they sharpen the research agenda, reveal changes we might have missed, or help us say “no” faster. We are already seeing clear value from our new tools.

TURNING PROCESS INTO INSIGHTS

In our Q3 commentary, we featured Google (Alphabet) as an AI stock. We remain convicted in Google’s positioning, but our work has made us increasingly enthusiastic about AWS, Amazon’s cloud infrastructure segment, as a beneficiary of where AI workflows are heading. In that spirit, our obsession is far more about the profit pools and platforms built on and leveraging AI than with the picks and shovels required to build out AI infrastructure. The market is focused on companies seeing a surge in sales as hyperscalers rush to build AI infrastructure, but we think that focus is misplaced.

When capex inevitably levels off at the hyperscalers (and it will), growth will evaporate and margins will compress at suppliers. Tier 2 suppliers in particular will see demand fall off a cliff, particularly as Tier 1 suppliers add capacity into a plateau in demand. As this generational buildout matures and growth capex tapers, investors will likely realize that an entire class of companies should never have had their earnings capitalized at such high multiples. Meanwhile, the companies building recurring revenues that will continue for decades receive little attention amidst the hype. The recurring revenues layer on slowly compared to the surge in orders for hot items, but the recurring revenues compound and do so at high incremental margins. Free cash flow is crimped as companies rush to meet growing demand; however, as growth slows, free cash flow will soar. These dynamics are opposite what today’s market obsessions will experience. Herein lies our obsession with the profit pools that are emerging today.

In that very same commentary we featured Google, we gave a brief shout-out to Amazon’s opportunity to thrive in building the “application layer” of AI by deploying smart orchestration across the retail business’ robotics and logistics layers. The orchestration opportunity will take time to play out, but meanwhile, we see AWS at a critical inflection point today. As discussed above, using the right model for the right task matters, and Amazon is uniquely positioned to benefit from a shift toward token efficiency and workflows built by AI but executed largely through non-AI processes. This has become increasingly clear in our own work and there is growing evidence that the most advanced companies deploying AI are moving this way themselves: durable value should accrue to platforms that can orchestrate models, data, compute, storage, security, and workflow execution at scale. Said differently, Amazon’s ability to remain model agnostic, while serving the lowest cost tokens, positions them uniquely to capitalize on AI adoption at scale.

AWS has been the platform that helped launch countless software, ecommerce and digital service companies and has empowered numerous older companies to migrate their digital infrastructure to the cloud. They have

done this with a combination of driving down the cost of compute, leveraging their proprietary chips and building an ecosystem of integrations around their offering.

Amazon was an early partner to Anthropic and owns a considerable equity stake in the company and more recently became an owner in OpenAI with an equity stake alongside a commitment from OpenAI to spend \$138 billion “to consume approximately 2 gigawatts of Trainium capacity through AWS infrastructure.”² Trainium is AWS’ custom AI chip, designed to compete with Nvidia’s GPUs at an industry-leading total cost of ownership. As CEO Andy Jassy explained, “Our Trainium2 chip has about 30% better price performance than comparable GPUs and is largely sold out. Trainium3, which just started shipping at the start of 2026 and is 30% to 40% more price performance than Trainium2, is nearly fully subscribed. And much of Trainium4, which is still about 18 months from broad availability, has already been reserved.”

At the heart of AWS’ AI offering is Amazon Bedrock. This is a hosted environment that can run many of the leading AI models and agents, as well as many of the open-source cost efficient ones. In a world where leading users of AI require a variety of models, alongside the ability to run non-AI programs, AWS is positioned to win because their scale is greater and their cost per unit (whether we’re talking tokens, CPU or memory) is lower than anyone else’s. Notably, Amazon is already seeing clear signs that the flywheel between AI workflows and core cloud infrastructure is starting to take hold and accelerate, as explained by Jassy:

And then at the same time, we're seeing very significant growth in our core business. And some of that are the migrations that have picked up from enterprises from on-premises to the cloud. But a lot of that is also as AI growth is exploding, it turns out that it leads to a lot of core growth as well, all the post-training, all the reinforcement learning, all the agentic actions and tool usage that these agents are using. And it fits with what you're asking about on the chip side, which is because we have an unusual collection of chips, we have the leading CPU chip in Graviton, and we have the leading price performance silicon AI chip in Trainium. It means that we're really unusually well positioned for the inflection that we're seeing and the type of growth that we're experiencing.

This growth is only just beginning and will accelerate as people move beyond experimenting with AI to running workflows built by AI, but the market has yet to recognize this reality. The opportunity in Amazon today feels similar to Google at this time last year. Due to AWS’ industry-leading scale during the rise of AI, growth rates are slower than other hyperscaler cloud peers; however, the absolute dollar volume of growth is incredible and accelerating today. This acceleration will continue throughout the year.

SAAS RISK TRACKER

We wanted to share one of our high value panels built by AI. This is a project we built in order to decipher the SaaS landscape and mine for opportunities where the market might be indiscriminately punishing software companies that are relatively inoculated from AI risk. The learnings have been actionable: since quarter-end, we have purchased two companies where this tracker helped us better understand the relevant risks. We will write about these purchases in our Q2 commentary. It has also kept us from acting on other companies that had been high up our watchlist.

Essentially what we have done is use a combination of quantitative and qualitative factors to assign a score that measures the risk a software company faces from AI. We defined the logic behind resilience and

2. <https://openai.com/index/amazon-partnership/>

identified a key of traits that would be strong indicators of resilience quantitatively. In our benchmarking, a low score is good, while a high score is bad. We have turned our scores into three separate indexes: a high, medium and low risk bucket, each of which we can track on their own. We can also track an aggregate index. Notably, although market performance is not an input in the model, actual market results have aligned strongly with the model's assessment of risk. While we are still tracking these data points prospectively—the backtested results and early tracking look promising.

We have overlaid fundamental data and given Claude the opportunity, knowing our worldview, to point us to mispriced market opportunities and to alert us to “value traps” where the fundamentals might appear compelling, but the risk is too great.

Further down the tracker, we have ranked and sorted every company in our SaaS universe based on the quality of their free cash flow. Each company is ranked objectively on free cash flow quality—high contribution from net income, low contribution from stock-based comp, little deferred revenue, etc. We also take note of the companies with the greatest improvements in free cash flow quality over time. We also analyze the composition of bookings. Companies with very short-duration bookings face different risks than those with longer-term contracts locked in.

The tracker mines the transcripts of each company and pulls out the most important quotes as it pertains to AI's impact on the business and ranks the quality of those insights. Companies who merely speak qualitatively receive less credit than those who quantify the benefits (and risks).

Last, we can click into any of the SaaS companies we track and see the key statements relating to AI displacement, renewal pricing, downsell/seat compression, build vs buy questions, profitability, renewal walls and AI monetization, amongst other factors. Everything is sourced and clickable back to the actual filings or transcripts. This has meaningfully accelerated our work in the SaaS space in a way that previously would not have been possible. We can cover more ground, get to “no” faster on certain companies, and develop a deeper appreciation for the persistence of certain businesses in ways that would have required a very different level of effort in the past.

The full quarterly snapshot of our tracker—covering every company in our SaaS universe along with their risk scores, free cash flow quality rankings, and AI-related management commentary—is available [here](#).

We believe this is an environment where disciplined active management matters. The opportunity set is changing, dispersion is meaningful, and the ability to separate durable fundamentals from temporary enthusiasm remains critical. We are excited about the opportunities in front of us and grateful for the trust you continue to place in us.

If anything in this commentary prompts questions, please reach out. You can contact any of us at 516-665-1945 or through our direct lines listed below.



Jason Gilbert, CPA/PFS, CFF, CGMA

Managing Partner, President

Direct: (516) 665-1940

jason@rgaia.com



Elliot Turner, CFA

Managing Partner, CIO

Direct: (516) 665-1942

elliott@rgaia.com



Ryan King

Partner

Direct: (516) 475-2475

ryan@rgaia.com