



PARNASSUS
INVESTMENTS®

THE PARNASSUS VIEW

MAY 11, 2026

Building on Strength: How Incumbents Win the AI Transformation

KEY TAKEAWAYS

- **AI is a durable trend** that is on the forefront of transforming industries worldwide. The investment in AI is historic in size, depth, breadth and imagination, but it will take time to see its full potential.
- **We believe most incumbents are well positioned to succeed.** Established companies with the deepest platform entrenchment and durable data are likely to strengthen their competitive moats by adding AI features.
- **The U.S. ecosystem offers advantages.** This self-reinforcing cycle of scale, capital access and global revenue is fueling U.S. leadership.
- **A “quality-first” approach is key.** Our investment theme favors companies with durable earnings that can compound through AI-driven cycles.

Every era of innovation reshapes the economy and its key players. Today, artificial intelligence (AI) is doing so at unprecedented speed, transforming how businesses operate, compete and create value. As investors, our job is to be able to spot the shifts, identify the leaders and invest into periods of growth.

The sheer magnitude of capital flowing into AI has grown at an exponential rate—funding everything from data centers to semiconductors to the infrastructure needed to power AI. It has ignited intense debate over whether we are witnessing a bubble or the start of a prolonged investment cycle.

Based on our decades of experience managing investments through multiple market cycles, we believe that AI is a durable trend. We are optimistic that it will bring about a long-term transformation that will unleash a huge wave of not only productivity enhancements, but also innovation in the economy. The emergence of agentic AI—automated AI systems that plan and execute multi-step tasks—has accelerated that transformation and shifted the central debate to which companies are best positioned to capture the most value.



Todd Ahlsten

Chief Investment Officer,
Portfolio Manager,
Parnassus Core Equity



Krishna Chintalapalli

Portfolio Manager,
Parnassus Value Equity



Andrew Choi

Portfolio Manager,
Parnassus Core Equity,
Parnassus Growth Equity

So far in 2026, the pullback in Magnificent 7 has been driven by increased scrutiny of the massive capex commitments by the largest hyperscalers. We do not see this as a failure to monetize AI, but rather the natural rhythm of capital-intensive technology cycles. Looking at past cycles for the internet, mobile and other innovations, it could take years or even decades of incremental progress for the innovation cycle to play out. Hyperscalers today face compute constraints that can lead to tradeoffs between their internal workloads or cloud customers. Without these constraints, the companies could see faster growth and higher return on invested capital (ROIC).

The AI transformation requires a new set of infrastructure, platforms and applications that can be deployed on a broad scale and can handle massive workloads. As companies figure out how to deploy AI for productivity and efficiency advantages, we expect that it will be a driver of continued U.S. exceptionalism, because companies in other countries won't be able to compete on the same scale. That's why we believe the U.S. incumbents building AI infrastructure, models and applications today represent compelling long-term investments—companies that have sturdy and durable cash flows and those that we think can compound value.

Because AI is changing so quickly, it's increasingly challenging to predict its evolution and assess its long-term impact on stock performance. While our view is that many large incumbents in the current AI value chain will likely remain good long-term investments, we also acknowledge that new startups outside our field of observation could disrupt the AI value chain and add to the uncertainty..

The Ecosystem Advantage

What sets American companies apart in AI is their ability to create self-reinforcing ecosystems that turn innovation into enduring competitive advantage. While competitors worldwide can achieve impressive technical breakthroughs with point solutions, it is difficult to create an ecosystem. Silicon Valley and American companies have established a virtuous cycle of hiring the world's best talent, securing significant sums of money for capital expenditures on infrastructure or development and building their business by monetizing and sustaining breakthroughs.

China's DeepSeek R1 model illustrates this dynamic. When it made headlines in early 2025 with claims of comparable performance at lower cost, it caused a brief selloff in AI stocks. But the excitement faded quickly, because DeepSeek lacked an ecosystem or commercialization engine that could help it scale adoption and keep pace with rapid advancements in technology. Meanwhile, Anthropic has secured strategic partnerships with **Microsoft** (MSFT), **Amazon** (AMZN) and **Alphabet** (GOOGL) for compute and distribution while building its own monetization flywheel through products like Claude Code and Cowork, reaching \$18 billion in annual recurring revenue in March 2026.

Unlike in the dot-com boom, today's leaders aren't starting from scratch—they're building on proven business models and cash flows from their anchor tenants. Alphabet has search, for example, while Microsoft has its Office suite. These companies can build infrastructure at massive scale, then monetize it across multiple use cases based on their existing business. For additional perspective on how the current situation compares to the dot-com boom and bust, review ["CIO Outlook—Will the Past Be Prologue for the AI Boom?"](#)

“

We believe incumbency remains a significant advantage, but the companies that will sustain it are the ones with data relationships, platform scale and workflow integration that make them indispensable in an agentic world.

”

In our view, AI providers that can effectively deploy agentic workloads to strengthen their businesses will be better positioned to succeed over unproven products and services from new providers trying to unseat them. The shift to agentic AI, however, can also create disruption risk for incumbents. We believe incumbency remains a significant advantage, but the companies that will sustain it are the ones with data relationships, platform scale and workflow integration that make them indispensable in an agentic world.

Maintaining the sheer scale required to operate globally leading business models demands setting and adhering to best practices in privacy and data security. Companies adopting AI applications in their operations will place a premium on providers that can handle their sensitive data and mission-critical business processes securely and ethically. Trustworthiness will be at least as much of a differentiator for AI model providers as the cost of delivery or technical superiority. These providers build trust by demonstrating commitment through large, ongoing investments and operating under robust regulatory frameworks.

Additionally, the ecosystems extend globally, which adds to the durability of the scale advantages. A significant portion of revenue at leading U.S technology companies comes from overseas, as shown in the table below. This established global base of customers should enable adoption of AI, because many of the U.S. incumbents' customers are already relying on their software and access to their data centers. They can unlock productivity and efficiency without installing an entirely new set of

EXHIBIT 1: Technology Ecosystems Extend Globally, Furthering Foundational Business Strength

Company	2025 International Revenue Share	Notes On Global Scale
Apple (AAPL)	57.1%	Apple’s revenue is diversified across global regions, with a large and critical share coming from Europe, Japan, China and other Asian markets.
Nvidia (NVDA)	53.0%	Nvidia’s AI chip sales are concentrated in hyperscale/data-center customers worldwide, with significant revenue in both the U.S. and Asia (notably Taiwan and China).
Alphabet (GOOGL)	52.0%	Revenue from its primary sources (Google Search, YouTube and Google Cloud) is nearly evenly split between the U.S. and international markets.
Microsoft (MSFT)	48.7%	Microsoft’s Azure cloud services and 365 enterprise software have nearly half their revenue sourced internationally.
Amazon (AMZN)	40.5%	While Amazon Web Services (AWS) is highly global, its North America segment (primarily its e-commerce platform) dominates overall sales. Overseas revenue reached \$300 billion in 2025.

Source: Source: Company filings. Data as of Dec. 25, 2025.

This U.S. ecosystem advantage manifests across three critical layers of the AI technology stack: 1) semiconductors; 2) hyperscalers that offer cloud infrastructure; and 3) enterprise software. Each subsector presents distinct opportunities shaped by different competitive dynamics.

Semiconductors: Formidable Scale and Barriers to Entry

The AI revolution requires enormous amounts of specialized computing power (also known as “compute”) to fuel the advanced processors and memory chips that enable machines to learn and generate responses at lightning speed. We are in a compute-constrained world. Demand for this computing capacity far outstrips supply, driven by an intense race to scale large language models and the rapidly growing use of these models. Put simply, companies building AI systems need more chips than the industry can produce. Another challenge data centers are facing is a bottleneck in electricity to power them. Given this demand, the semiconductor industry is on track to grow from roughly \$631 billion in 2024 to more than \$1 trillion by the end of this year.¹

The current semiconductor cycle has fundamentally changed how the industry operates, splitting it into two distinct cycles: AI and non-AI. Non-AI cycles relate to mainstream automotive, industrial and PC applications. Traditional semiconductor markets serving the non-AI businesses have remained relatively weak, while AI-related demand has expanded, particularly in graphics processing units (GPUs). This creates demand for custom application-specific chips (ASICs) from chipmakers like **Broadcom** (AVGO), central processing units from **AMD** (AMD) and advanced memory from **Micron** (MU), among others.

While high-performance GPUs remain essential for AI adoption, there is a broadening of compute requirements as AI moves from pure model-training to inferencing and real-world applications. It’s a race to the finish, where we’re seeing companies selecting the best-of-breed chips rather than taking the time to conduct tests and try out different types of chips. The market is moving toward adoption at a very rapid pace.

We believe the physical constraints of the buildout in 2026 are actively preventing compute capacity from overshooting demand. Power availability, memory supply and financing constraints are all limiting the pace at which compute can actually be installed. Because demand, driven by continuous workloads of agentic AI, is rising faster than these constraints allow supply to respond, we believe the semiconductor cycle is more likely to be elongated than to overshoot. We continue to monitor LLM scaling limits, hyperscaler capital discipline and enterprise adoption rates as the key indicators that would signal a change in this view.

One emerging risk we’re watching is that the exponential ramp-up of data center investments is starting to trim the free cash flows of even the largest hyperscalers. For now, operating cash flows at most of the hyperscalers exceed capital spending, as the chart below shows. In some cases, however, they are starting to issue debt and use more innovative ways to finance these large projects. The shift to debt financing is making the cycle more interest-rate sensitive than it was in its early stages. This is a sign that the demand growth from the AI data center buildout may be approaching a funding speed limit. At the same time, emerging shortages of power and semiconductors, such as memory chips, also reveals the physical limitations of keeping up with AI digital compute needs.


“

We are in a compute-constrained world. Demand for this computing capacity far outstrips supply, driven by an intense race to scale large language models and the rapidly growing use of these models.

”

EXHIBIT 2: HYPERSCALER CASH GENERATION MOSTLY EXCEEDS SPENDING

Cash flows from operations vs. capital expenditures, 2026 estimates



Source: FactSet, Parnassus Investments.

Hyperscalers: The Infrastructure Enablers


Hyperscalers operate massive data centers that provide scalable technical infrastructure as a service, a necessary building block for digital products and applications that businesses may create. Leading hyperscalers include the “Big 3” of Microsoft, Alphabet and Amazon, as well as more specialized cloud infrastructure providers like **Oracle** (ORCL). Hyperscalers combine several critical advantages: broad scale, access to GPUs, strong cash flows that can finance capital expenditures and access to data for training AI models.

The hyperscaler business model benefits from powerful economies of scale that are difficult to replicate. The core of the cloud business is shared infrastructure. The more people use it, the more leverage you can get out of the system. This dynamic is accelerated by AI, which is driving companies to complete the migration from managing on-premise systems, which require more resources from the company to maintain, to cloud platforms, which are maintained mostly by the hyperscaler. With the cloud model, the more clients the hyperscalers can add to the cloud, the more payback they can get out of the system. On top of that, as companies make the transition to cloud-based software applications, it becomes easier for the company to explore the use of the hyperscaler’s AI tools, making AI a significant driver for overall cloud adoption.

In recent months, investors have grown skeptical about whether hyperscalers are overinvesting in AI, especially if costly GPUs depreciate quickly as more powerful chips emerge. The chart in Exhibit 3 shows how major hyperscalers are planning to spend nearly \$670 billion on AI data centers in 2026, a sevenfold increase since 2020.

EXHIBIT 3: AI ARMS RACE: CAPEX SURGING AMONG HYPERSCALERS

Annual capital expenditures by company



Source: FactSet, Parnassus Investments. 2026 figures are estimates.

We believe these concerns overlook the practicality of these investments. The chips and infrastructure deployed can be used for other purposes, such as for accelerating existing workloads, AI use cases beyond merely large language models (LLMs) or other emerging use cases that demand compute. Infrastructure investments aren't likely to become obsolete, although there are nuances in terms of where the greatest opportunities may be. Moreover, as compute costs decline over time, more and more companies should continue to adopt AI, which should help to grow the addressable market for the hyperscalers' services.

The debate continues. We believe hyperscalers are well positioned to see that ROI over the long term. In our view, the current AI infrastructure buildout is developing similarly to the early stages of the cloud computing buildout for these same hyperscalers. When Amazon first committed billions of dollars to cloud-based data centers, investors similarly questioned whether the spending would generate suitable returns. Instead, declining unit costs and scale benefits unlocked entirely new categories of demand and returns on data center spending have been well above required thresholds. We view the AI-based data center buildout as potentially following a similar path, requiring similar levels of patience upfront, but confident in the long-term durable advantages these companies are building as AI usage proliferates the global economy.

There is also the risk that a short-term period of overcapacity could trigger a mini crash. While we could see some short-term losses, it would effectively cut the price for compute. This could then attract a new set of companies to take on the infrastructure. After the dot-com bust, the internet didn't cease to exist. Similarly, such an event wouldn't represent a fundamental failure of the technology, but rather the natural rhythm of capital-intensive technology cycles we've observed repeatedly in PCs, telecommunications and cloud computing.

Software: Selectivity During Uncertainty

Much attention has focused on infrastructure spending, while software companies appear to be left for dead. The companies' outlooks have become the subject of a debate. Will they become irrelevant? We don't think the death of enterprise software is imminent, but we do think the shift to agentic AI has made the investment case significantly more selective.

Not all incumbents will be equally protected by agentic AI. The companies that present opportunities share specific characteristics. They are embedded at the platform level, not merely the application level; their data relationships create switching costs that autonomous AI systems cannot easily replicate; and they have credible roadmaps for deploying agentic capabilities that expand their value to customers rather than exposing them to substitution. Software businesses with applications that agentic AI can more readily replace face a more uncertain competitive outlook.

Enterprise software companies have been integrating AI features into their solutions, leading to a prolonged phase of exploration and experimentation while the companies figure out how to seamlessly embed and monetize AI capabilities. We see AI agents as an opportunity for software providers with deeply embedded platforms. They can serve as an intelligence layer that makes existing software both more powerful and more difficult to displace. We are focusing our software holdings on companies with data depth and platform incumbency that can benefit from agentic workloads. We view the software providers with entrenched systems of record as being best positioned to monetize AI through enhanced user experience and agents. This includes Microsoft and vertical software providers serving specific industries such as **Guidewire** (GWRE) in insurance and **Appfolio** (APPF) in real estate.

We believe these companies can further customize the software for a customer's needs and potentially expand their addressable market, creating further opportunities for growth. A key reason software companies haven't grown as fast as other AI infrastructure providers since the release of ChatGPT is the continuing digestion of the surge in demand and investment for software as a service during the COVID-19 pandemic. We believe AI-enabled features could revitalize and expand their potential as new solutions and pricing models evolve to better capture the value proposition.

A Quality-First Investment Philosophy

We are only in the early stages of seeing this rapidly evolving transformation take shape, so we write this view with a sense of humility. Our job as investors is to create value for our shareholders. We do that by owning stocks with an adequate return potential in relation to risk. As capital-intensive innovation cycles play out, this means identifying and owning stocks that aren't yet pricing in their ultimate earnings growth, their ROIC and their cash flows.

Applying this investment approach to AI means weighing the transformative long-term potential of the new technology against its short-term risks. We align with the vision of tech incumbents that AI represents not just a technological shift, but a transformational wave in how businesses operate and compete in the coming decades. The companies building the infrastructure, platforms and applications that enable that transformation are positioned to capture and create substantial value over the long term. ■

“

We don't think the death of enterprise software is imminent, but we do think the shift to agentic AI has made the investment case significantly more selective.

”

RESPONSIBLE AI AND LONG-TERM VALUE CREATION

At Parnassus, we take investment stewardship seriously, recognizing our responsibility as active owners through proxy voting and company engagements. We draw upon our 40 years of sustainable investing leadership experience to focus on the risks and opportunities to preserve and grow shareholder value.

AI represents both tremendous opportunity and significant risk, especially with unintended consequences for customers. It introduces governance challenges that demand rigorous oversight. Our stewardship efforts concentrate on ethical AI use, transparency and board-level oversight with the goal of protecting people and preserving trust.

With regard to a company's AI usage, we monitor the governance, policies and practices for:

- **Bias and fairness.** AI algorithms can drive consequential decisions in hiring, lending and healthcare, among others. Discrimination or unfairness in these sensitive domains threatens corporate reputations and stakeholder relationships.
- **Privacy and security.** As AI systems train on vast datasets, we aim to ensure that companies respect individual rights, copyright and intellectual property protections, especially regarding product development.
- **Transparency.** We emphasize clear disclosure when customers interact with AI systems and clarity around how decisions are made. This is increasingly vital as AI blurs the line between human and machine judgment.

We encourage portfolio companies to align with the National Institute of Standards and Technology (NIST) AI Risk Management Framework, which is the closest thing we have to a U.S. gold standard for AI governance. Through detailed policy review and alignment with best practices, we aim to help companies protect their business while navigating AI's transformative potential. It's an incredibly important time for stewards of capital to be paying attention to the approaches companies are taking to ensure the best possible outcomes for customers, the company and its stakeholders.

To learn more, visit www.parnassus.com.

¹ Semiconductor Industry Association, "Global Semiconductor Sales Increase 29.8% Year-to-Year in November," <https://www.semiconductors.org/global-semiconductor-sales-increase-29-8-year-to-year-in-november/>. Jan. 8, 2026.

PIL-931792-2026-05-05

Magnificent Seven (Mag 7): A group of seven large, dominant technology companies whose performance significantly influences the stock market. The companies are Apple, Microsoft, Amazon, Alphabet (Google), Meta Platforms (Facebook), Tesla, and Nvidia.

Cash Flow: The net amount of cash and cash equivalents being transferred in and out of a company.

Return on Invested Capital (ROIC): A measure of a company's efficiency in allocating capital to profitable investments. It is calculated by dividing net operating profit after tax by invested capital.

S&P 500® Index: (registered trademark of The McGraw-Hill Companies, Inc.) is an unmanaged index of 500 common stocks primarily traded on the New York Stock Exchange, weighted by market capitalization. Index performance includes the reinvestment of dividends and capital gains. An individual cannot invest directly in an index. An index reflects no deductions for fees, expenses or taxes, but mutual fund returns do.

Performance may be affected by a smaller number of issuers than if it were a diversified, less concentrated fund. Further, the Fund may experience greater losses as a result of a single issuer's unfavorable market or economic conditions or other adverse developments impacting the market value of the issuer's securities.

ENVIRONMENTAL, SOCIAL AND GOVERNANCE (ESG) GUIDELINES: The Fund evaluates ESG factors as part of the investment decision-making process, considering a range of impacts they may have on future revenues, expenses, assets, liabilities and overall risk. The Fund also utilizes active ownership to encourage more sustainable business policies and practices and greater ESG transparency. Active ownership strategies include proxy voting, dialogue with company management and sponsorship of shareholder resolutions, and public policy advocacy. There are no assurances the Funds will meet their investment objectives or that their ESG strategies will be successful.

Mutual fund investing involves risk, and loss of principal is possible. **Past Performance cannot predict future results.**

ETFs are subject to additional risks that do not apply to conventional mutual funds, including the risks that the market price of an ETF's shares may trade at a premium or discount to its net asset value, an active secondary trading market may not develop or be maintained, or trading may be halted by the exchange in which they trade, which may impact an ETF's ability to sell its shares. Shares of any ETF are bought and sold at market price (not NAV) and are not individually redeemed from the ETF. Brokerage commissions will reduce returns.

For the current holdings of the [Parnassus Core Equity Fund](#), the [Parnassus Growth Equity Fund](#), the [Parnassus Value Equity Fund](#), the [Parnassus Mid Cap Fund](#), the [Parnassus Mid Cap Growth Fund](#), the [Parnassus International Equity Fund](#), the [Parnassus Core Select ETF](#), and the [Parnassus Value Select ETF](#), please visit each fund's individual holdings page. Fund holdings are subject to change at any time.

©2026 Parnassus Investments, LLC. PARNASSUS, PARNASSUS INVESTMENTS and PARNASSUS FUNDS are federally registered trademarks of Parnassus Investments, LLC.

The Parnassus Funds are distributed by Parnassus Funds Distributor, LLC.

Before investing, an investor should carefully consider the investment objectives, risks, charges and expenses of a fund and should carefully read the prospectus or summary prospectus, which contain this and other information. The prospectus or summary prospectus can be found on the website, www.parnassus.com, or by calling (800) 999-3505 for a mutual fund prospectus or (855) 514-4443 for an ETF prospectus.