

Phase 4 AI: The Adopter's Advantage

At MoatPeak, we believe we are living through a rare kind of market moment, the kind that only appears a few times in an investing lifetime. (see *Figure 1*)





The Next Chapter in AI: From Silicon to Profitability

Why the most important investment opportunity of 2026 is not in technology creation, but in its adoption.

We believe we are at a critical inflection point. For the past three years, the market has focused on the creators of AI infrastructure, like \$NVDA.

Now, the capital cycle is turning. The primary value capture is migrating from the **builders of AI** to the **users of AI**.

This presentation outlines our framework for identifying the leaders of this new phase and navigating its unique risks.

Figure 1. The next chapter in AI: from silicon to profitability

Key takeaway: We see a shift in value capture from AI builders to the companies that use AI to improve margins and free cash flow.

For years, the center of the AI story was clear: it was about the builders. Investors chased chipmakers, cloud providers, and the companies that supplied the raw “silicon and chips” powering the AI revolution. Owning the infrastructure was the main way to “own AI.” (see Figure 2)

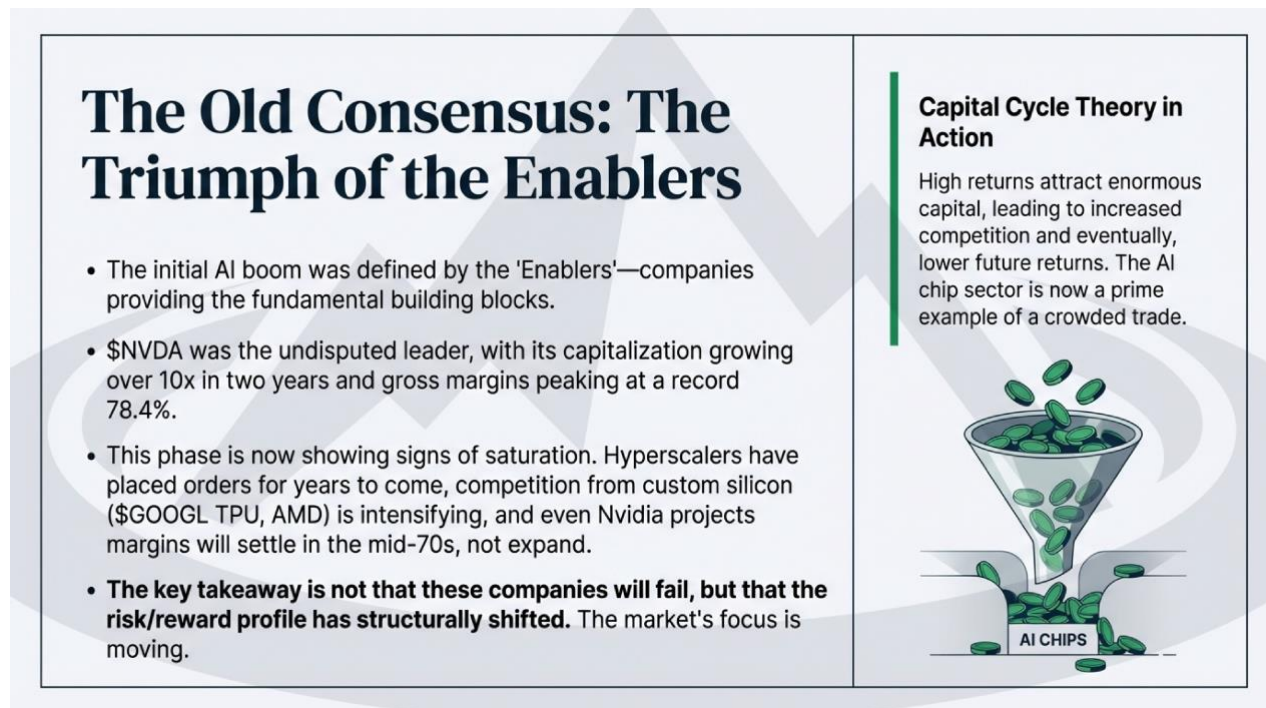


Figure 2. The old consensus: triumph of the AI enablers

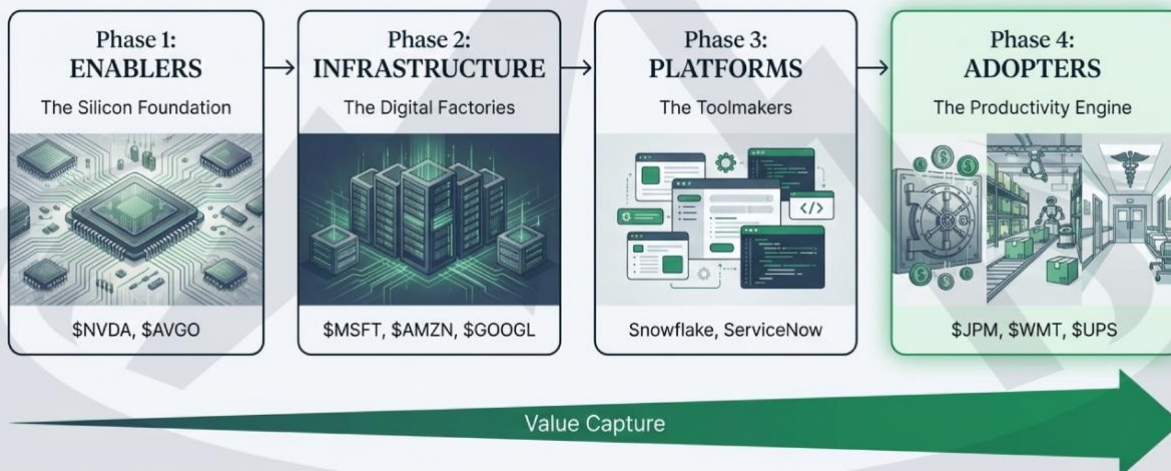
Key takeaway: Phase 1 rewarded the infrastructure builders; in 2026 the risk/reward is shifting toward adopters rather than track-layers.

That is now changing.

In 2026, the more important question is not who builds AI, but who uses it best. The next wave of value, in our view, will come from the traditional companies that quietly integrate AI into their day-to-day operations and turn it into real productivity, higher margins, and stronger free cash flow.

Goldman Sachs describes this shift as a move from Phase 1, the era of “silicon and chips,” to Phase 4, the era of “balance sheets and cash flow.” A simple metaphor helps: Phase 1 was about laying the railroad tracks of the AI economy. Phase 4 is about the new railroad magnates, the operators who use those tracks to build entire business empires. (see Figure 3)

A Framework for a Revolution: The Four Phases of AI



To understand where the market is going, we must first have a map of where it has been. We view the AI revolution not as a single event, but as a four-phase value chain. While the market has been fixated on Phases 1 and 2, we believe the most significant, and least appreciated, opportunity now lies in Phase 4.

Figure 3. A framework for the AI cycle: the four phases

Key takeaway: AI’s value chain is best understood as a four-phase cycle, with the most underappreciated opportunity in Phase 4 adopters.

The numbers behind this shift are substantial. Goldman estimates that AI-driven efficiencies could add around +40 basis points to S&P 500 net margins in 2026 and another +60 basis points in 2027. That might sound abstract, but it is not: those extra basis points could support 13–15% annual earnings-per-share growth for the broad market. In other words, AI is not just a story about future potential; it is a very real engine for earnings. (see Figure 4)

The Thesis: AI's Productivity Dividend is Coming

Goldman Sachs projects AI-driven efficiencies will add +40 to +60 basis points to S&P 500 net margins by 2026-2027.



Figure 4. AI productivity dividend: the margin expansion thesis

Key takeaway: Goldman’s margin math implies AI efficiency can meaningfully lift S&P 500 profitability and support double-digit EPS growth.

We are already seeing what this looks like in practice.

Citigroup now saves roughly 100,000 developer hours every week thanks to AI tools that automate routine coding and maintenance work. JPMorgan has used its COiN platform and similar systems to remove about 360,000 hours of legal work per year, freeing high-cost professionals to focus on higher-value tasks. Walmart has quietly built one of the most advanced AI operations in retail and achieved a 100-fold productivity improvement in its product cataloging process. Behind each of these statistics are thousands of everyday decisions, forms, tickets, and tasks that are now handled faster, cheaper, and more accurately than before. *(see Figure 5)*

From Theory to Reality: How Adopters Are Creating Value

Banking & Finance



Focus: Automating knowledge work and controlling headcount.

- **\$JPM**: Saves **360,000 hours** of legal work annually with its COiN platform.
- **\$C**: Saves **100,000 developer hours per week** through automated code review.
- **BNY Mellon**: Has **117 distinct AI solutions** already in production, from lead generation to payment processing.

Retail & Warehousing



Focus: Optimizing supply chains and back-office functions.

- **\$WMT**: Created four proprietary “super-agents” for customers, employees, and suppliers. Accelerated product card processing **by 100x**.
- **High-End Retail**: A key sub-theme is the outperformance of retailers targeting high-income consumers (\$100k+), who have shown more resilient spending. These companies (\$Nordstrom, \$Williams-Sonoma) can better absorb AI investment costs and pass them on.

Figure 5. From theory to reality: early evidence of adopter-led gains

Key takeaway: Real-world deployments at large incumbents already point to tangible labor savings and operational leverage from AI adoption.

To formalize this theme, Goldman created the AI Productivity basket (\$GSXUPROD). What makes it interesting is what it excludes: it deliberately leaves out the obvious tech and chip names. Instead, it focuses on adopters – the banks, retailers, industrials, healthcare companies, and logistics firms that are actually using AI to run their businesses more efficiently. (see Figure 6)

Identifying the Leaders: Inside the Goldman Sachs Productivity Basket (GSXUPROD)

GSXUPROD underperformed the S&P 500 in 2023-2024. We see this not as a failure, but as proof that the market has **not yet priced in the Phase 4 productivity gains**. This is why Goldman calls it “the most important trade for 2026.”

What is GSXUPROD?

- A curated portfolio of ~50 **non-IT companies** poised to be the biggest beneficiaries of AI adoption.
- Selection is based on tangible criteria: high labor costs, significant automation potential, and explicit management commentary on AI initiatives.

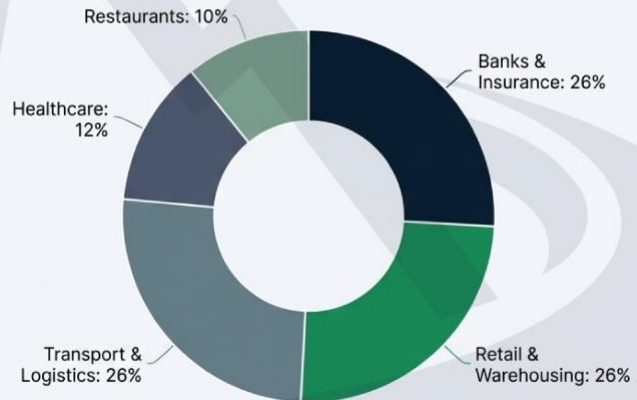


Figure 6. Inside GSXUPROD: the AI adopters basket

Key takeaway: GSXUPROD is positioned as a curated basket of non-tech beneficiaries of AI adoption—still underpriced in our view given its underperformance.

So far, this basket has underperformed the broader market. That underperformance is not a sign that the idea is wrong; in our view, it is a sign that the market has not fully woken up to the scale of the change. It suggests that many investors are still anchored in the Phase 1 mindset, chasing the enablers and overlooking the adopters.

None of this, of course, comes without risk. We are closely watching three big pressure points. First, there is the organizational problem of “pilot purgatory,” where companies run endless AI experiments that never scale into real, profit-generating systems. Second, there are physical bottlenecks in energy, grid capacity, and essential hardware. Third, there is mounting credit risk in companies that are using heavy leverage to fund their AI build-outs. (see Figure 7)

The AI Value Chain: Risk and Reward at Each Stage

Phase	Dominant Players	Core Thesis	Key Risk Profile
1. Enablers	\$NVDA, \$AVGO	Selling the "shovels" in a gold rush	Saturation & Margin Compression: Crowded trade, rising competition.
2. Infrastructure	\$MSFT, \$AMZN, \$GOOGL, \$ORCL	Massive CapEx to build AI-ready clouds	Capital Intensity & Debt Load: Requires huge investment before monetization.
3. Platforms	Snowflake, Databricks	Building software to enable AI adoption	Execution Risk: High R&D spend, path to profitability uncertain for many.
4. Adopters	\$JPM, \$WMT, \$JBHT	Using AI to drive margin and FCF	Implementation Failure: "Pilot Purgatory," organizational resistance.

The investment thesis evolves with each phase. While earlier phases focused on revenue growth and infrastructure build-out, Phase 4 is a direct play on profitability and operational efficiency. This is the shift from promises to proven ROI.

Figure 7. AI value chain: risk and reward at each stage

Key takeaway: The Phase 4 rotation is not risk-free: we monitor implementation failure ('pilot purgatory'), physical constraints, and credit tail risks.

Even after accounting for these risks, we believe the rotation into high-quality AI adopters is one of the most significant, and still underappreciated, investment themes for 2026. The rest of this report explains why.

2.0 The Structural Impact on U.S. Markets

If you want to navigate the next market cycle, watching daily price swings is not enough. You need to understand how AI is reshaping the foundations of both equity and credit markets.

AI is no longer a "feature" or a niche technology. It is becoming a general-purpose technology, similar in importance to electricity or the internet. That means it affects how companies earn profits, how investors value those profits, and how capital gets priced and allocated.

U.S. Equities

The most visible effect in equities will come through profitability.

With AI, a company can grow revenue without growing headcount at the same pace, or it can hold prices steady while reducing costs behind the scenes. When this happens across enough companies, it

shows up as margin expansion at the index level. This is what Goldman is pointing to with the projected +40 basis points of S&P 500 net margin expansion in 2026 and +60 basis points in 2027.

The market is already beginning to differentiate between companies that engage with AI and those that do not. So far this year, companies that mention “AI” on their earnings calls have seen shares rise by around +13.9%, while those that make no mention of AI are up about +5.7%. Today, some of that gap is still driven by excitement and narrative. Over the coming years, we expect the gap to be driven more and more by actual improvements in productivity, revenue per employee, and margins.

In other words, Phase 4 is when AI stops being something companies talk about and becomes something that shows up in the numbers.

U.S. Fixed Income and Credit

On the fixed income side, the AI story looks different but is just as important.

Building AI infrastructure is expensive. Data centers, specialized chips, cooling systems, fiber networks, and the power to run all of it require massive capital spending. Many companies are choosing to finance these projects with debt rather than equity, which shifts part of the AI story into the credit markets.

Oracle (\$ORCL) is a telling example. To fund its AI ambitions, Oracle has raised about \$18 billion in bonds, taken on \$18 billion in project financing, and put in place a \$38 billion credit line. The result is a large and complex debt stack. Investors have noticed. Oracle’s credit default swap (CDS) spreads, which reflect the cost of insuring its debt against default, have risen into the 126–139 basis point range. That is the highest level since the 2008 financial crisis. (*see Figure 8*)


The Canary in the Coal Mine: AI’s ‘Credit Tail Risk’

Case Study: Oracle (\$ORCL)

- Oracle is aggressively investing billions in AI data centers, accumulating massive debt to fund its infrastructure race.
- The company has taken on ~\$18B in bonds, \$18B in project financing, and has a \$38B credit line.

The Market’s Verdict

- Credit Default Swap (CDS) spreads on Oracle—the cost to insure against its default—have soared to 126-139 bps, the highest levels since the 2008 financial crisis.
- Morgan Stanley has called it the “preferred vehicle to hedge AI risk.”



Broader Implication: Oracle is not an isolated case. If the promised ROI from these massive investments fails to materialize quickly, or if interest rates rise, a wave of credit stress could hit the market, with high-yield debt (\$HYG) being particularly vulnerable.

Figure 8. Credit tail risk: Oracle as an early warning signal

Key takeaway: The debt-funded AI build-out is showing up in credit markets; Oracle’s CDS spread move is a key stress indicator to monitor.

Because of this, Oracle has become a kind of “credit canary” for the AI build-out. Institutional investors increasingly use it as a way to hedge against an AI-related credit shock. If stress like this remains contained in a few large names, the market can absorb it. But if it spreads, it could spill over into high-yield bond ETFs such as \$HYG, signaling that the AI investment cycle is putting real pressure on the broader debt markets.

These dynamics are not confined to the United States. The need for power, hardware, and materials is global, and the ripple effects are being felt worldwide.

3.0 The Ripple Effect on European and Global Markets

Although the most visible AI adoption stories come from U.S. companies, the physical foundations of this revolution – energy, grids, transformers, metals, and specialized equipment – are inherently global.

Consider something as mundane, yet essential, as a large power transformer. These devices step up and step down voltage on grids and are critical to connecting new data centers and industrial facilities. Today, the world is facing long waiting times for them. Lead times for large power transformers have stretched to more than two years in many regions. A data center project in Virginia and a similar

project outside Frankfurt may be competing for essentially the same scarce global manufacturing capacity.

The corporate response tells its own story. Hitachi Energy has announced a \$1.5 billion global investment plan to expand its transformer and grid equipment manufacturing. Siemens is committing around \$150 million to expand a key transformer plant in the United States. These are not tech companies in the traditional sense, but they are becoming critical suppliers to the AI era.

At the same time, AI's hunger for power is reshaping commodity markets. Large data centers consume huge amounts of electricity, which supports long-term interest in uranium and other energy-related themes. The need to strengthen and extend power grids increases demand for industrial metals, especially copper, where companies like Freeport-McMoRan (\$FCX) play a pivotal role.

All of this means that European and Asian industrial and energy companies are not just bystanders to the AI boom. They are essential suppliers. Their order books, capacity decisions, and pricing power will be influenced by AI just as surely as any Silicon Valley software firm.

4.0 Sector Deep Dive: The Engines of Phase 4

To see where AI adoption may create the most value, it helps to look more closely at the sectors inside the Goldman Sachs AI Productivity basket (\$GSXUPROD). This basket is not full of glamorous tech names. Instead, it is overweight in the kinds of companies that have lots of people, lots of process, and lots of repetition. Those are precisely the places where AI can have the biggest impact.

4.1 Financials: Banks and Insurance (about 26% of the basket)

Financial institutions are, at their core, information-processing machines. They take in data, assess risk, apply rules, and make decisions. That makes them perfect laboratories for AI.

At JPMorgan, for example, the COiN platform reads and interprets complex legal contracts that once required hours of attention from human lawyers. It is estimated to save 360,000 hours of legal work every year. Those are 360,000 hours that can be redirected from low-value document review to higher-value activities.

Bank of America has rolled out its virtual assistant Erica to tens of millions of customers. Routine questions that once led to long phone calls or in-branch visits can now be handled in seconds through a conversational interface. Citigroup is deploying internal AI tools across its 180,000-person workforce to streamline tasks ranging from IT support to internal reporting.

In insurance, AIG has launched "AIG Assist," an AI-driven tool that accelerates underwriting and improves claims handling. Instead of waiting days for a manual review, many policy decisions and claim assessments can now be made in far less time, with greater consistency.

In all of these cases, AI is not a side project. It is being integrated into the main arteries of the business.

4.2 Consumer and Logistics: Retail and Warehousing (about 26%)

Retail is both a customer story and a logistics story. On the front end, buyers expect personalization, speed, and convenience. On the back end, retailers manage enormous catalogs, complex supply chains, and thin margins. AI touches all of this.

Walmart (\$WMT) is a good illustration. It has developed four major internal AI “super-agents” that support customers, employees, suppliers, and developers. Perhaps the most striking achievement is a 100-fold increase in productivity in its product cataloging process. What used to involve slow, manual entry and checking can now happen quickly and with far fewer errors.

In the private markets, Lineage Logistics has used proprietary AI systems to coordinate its global network of temperature-controlled warehouses. The result has been double-digit percentage productivity gains across the network. Trucks are dispatched more efficiently, storage is better utilized, and spoilage is reduced.

A pattern is also emerging within retail itself. Companies that serve more affluent customers, such as Nordstrom or Williams-Sonoma, have often seen more stable revenue and higher spending even through economic uncertainty. That financial resilience gives them more room to invest in technology, including AI, without sacrificing short-term performance. Over time, this can compound into a meaningful competitive edge over lower-margin peers that cannot invest as aggressively.

4.3 Industrials: Transport and Logistics (about 26%)

Behind the scenes of the global economy are fleets of trucks, containers, warehouses, and distribution centers. Coordinating all of this is a classic “hard problem”: many moving parts, thin margins, and huge rewards for getting it right.

AI is becoming the invisible dispatcher in this system. At J.B. Hunt (\$JBHT), AI tools have automated around 60% of carrier verifications, 73% of order intake, and 80% of invoicing. Each of these tasks once involved significant manual effort. Automating them not only saves time and labor costs, it also reduces errors, accelerates billing, and shortens the cash conversion cycle.

XPO Logistics (\$XPO) is using AI to optimize its routes and pricing decisions. By better matching trucks, loads, and routes, it has achieved a 12% reduction in empty miles. Those are miles where a truck is on the road but not earning revenue. Cutting them by 12% directly improves fuel efficiency, asset utilization, and profitability.

4.4 Healthcare (about 12%)

Healthcare combines high stakes with heavy bureaucracy. Doctors and nurses are often surrounded by forms, codes, billing rules, and compliance requirements. While AI has exciting clinical potential, the earliest and most reliable returns are showing up in administrative workflows.

Networks such as HCA Healthcare (\$HCA) are using AI to automate revenue cycle management, from coding and billing to follow-ups and appeals. Elevance Health (\$ELV) and other insurers are deploying AI to process claims more quickly and accurately, reducing denials and rework. The immediate benefits are fewer mistakes, faster payments, and lower overhead costs, all of which support margins without compromising patient care.

4.5 Consumer Discretionary: Restaurants (about 10%)

Restaurants live on thin margins and face constant pressure from rising wages and tight labor markets. Here, AI and robotics are starting to reshape work in very tangible ways.

Chipotle (\$CMG) is testing a robot called “Autocado” to handle the repetitive work of preparing avocados for guacamole. Wendy’s is piloting “FreshAI,” a voice-based AI system that takes drive-through orders with consistent accuracy and speed. Sweetgreen has gone further, reporting a 700 basis point reduction in labor costs in its robotic-enhanced kitchens and higher average check sizes thanks to AI-driven recommendations at self-service kiosks.

Taken together, these examples show how AI is moving from PowerPoint slides into kitchens, billing departments, call centers, and warehouses. (see Figure 9)




Figure 9. AI adoption is sector-agnostic

Key takeaway: Phase 4 benefits are emerging across logistics, healthcare, and consumer services—not just traditional technology sectors.

4.6 The Second-Order Play: The Infrastructure Behind It All

There is a second layer to the Phase 4 story that is easy to overlook: the companies that supply the essential components to make all this possible.

In the railroad era, fortunes were not only made by the rail barons, but also by those who supplied the rails, ties, and signaling equipment. In today’s AI economy, those suppliers include companies like Vertiv (\$VRT) and Eaton (\$ETN), which provide the power and cooling systems that keep data centers running. They also include producers of uranium (such as \$URNM) and copper (such as \$FCX), whose materials are crucial for generating and transmitting the electricity that AI demands. *(see Figure 10)*



Figure 10. Second-order beneficiaries: investing in 'second-line shovels'

Key takeaway: Beyond adopters, suppliers of power, cooling, grid equipment, and essential materials can capture structural AI-driven demand.

These “second-order” businesses may not be the ones showcasing chatbots or generative models, but they stand to benefit from the structural increase in demand created by AI adoption.

5.0 Grey Rhinos: The Structural Bottlenecks

When markets think about risk, they often obsess over rare, surprising disasters – the famous “Black Swans.” We believe the more important risks to the AI thesis are the ones that are obvious, slow-moving, and still widely underestimated. We call these “Grey Rhinos.” *(see Figure 11)*

The Mentor’s View: Acknowledging the New Bottlenecks

MoatPeak Mental Models in Focus

- **"Inversion":** "To understand who will succeed in Phase 4, we must first understand the primary reasons for failure."
- **"Second-Order Thinking":** "The AI boom has obvious first-order effects (demand for chips). We must ask, 'And then what?' to uncover the hidden constraints."



We see three structural bottlenecks that could slow the AI transition:

Figure 11. Acknowledging the new bottlenecks

Key takeaway: We frame the main risks as 'grey rhinos'—obvious structural bottlenecks that can slow the shift from AI slides to cash flow.

Three of them stand out.

The first is energy. By 2028, the United States is expected to face a power shortfall of roughly 45 gigawatts. At the same time, every additional gigawatt of AI computing capacity is estimated to require around 50 to 60 billion dollars in grid investment. If the power is not available when and where it is needed, data centers will be delayed or scaled back. The companies that already control access to reliable power will hold a valuable advantage.

The second is the grid and transformer system. Large power transformers are already in tight supply. Forecasts point to a 30 to 40 percent deficit in available units by 2027. Lead times have stretched to about 115 to 130 weeks, and prices have risen 60 to 80 percent. That combination of scarcity, delay, and higher cost can easily push new projects back by six to twelve months and slow the payoff from AI-related capital spending.

The third is labor. AI may automate cognitive tasks, but building the infrastructure still requires people with very specific skills. The United States is estimated to face a shortfall of 63,000 to 140,000 electricians over the coming years. Microsoft’s ambitious “Stargate” data center project alone is expected to require 12.6 million man-hours. When labor is scarce, projects move slower and become more expensive, which reduces the internal rate of return for investors. *(see Figure 12)*

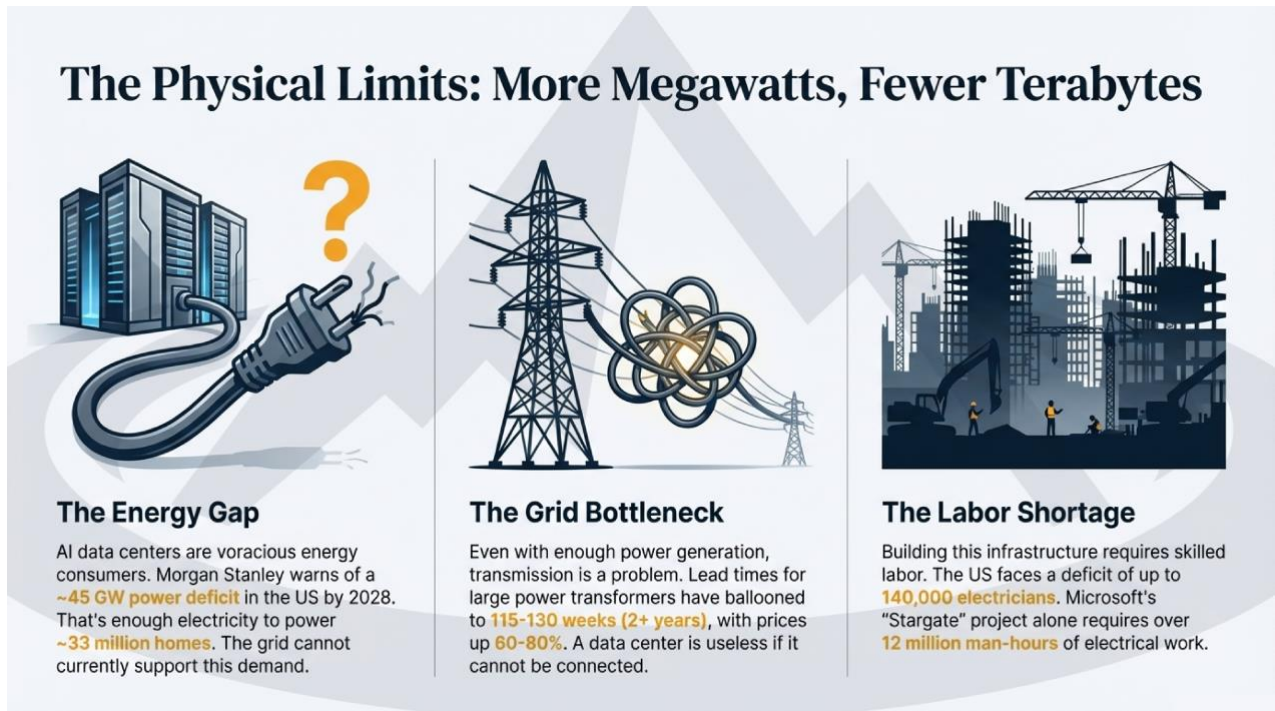


Figure 12. The physical limits: energy, grid, and labor constraints

Key takeaway: AI's infrastructure needs face binding constraints in power availability, transformer lead times, and skilled labor supply.

The Organizational Bottleneck: "Pilot Purgatory"

Beyond these physical constraints lies a quieter, but equally dangerous, bottleneck: the way organizations themselves behave.

McKinsey's research suggests that about 67 percent of companies are stuck in AI pilot phases and that only 39 percent are seeing meaningful gains in EBIT from their AI efforts. Many companies know they "need to do something with AI," but struggle to turn that intention into scaled, sustainable systems.

This "pilot purgatory" usually stems from a mix of factors: weak or fragmented leadership, poor-quality data scattered across systems, and cultural resistance from employees who fear change or do not trust the tools. The result is a lot of interesting experiments that look impressive in slide decks, but never quite make it into the real workflows that drive profit.

Over time, this will create a sharp divide between those who learn to execute and those who do not. The market will eventually reward the companies that move beyond pilots and embed AI deeply into their operations.

6.0 Our Scenarios and Probabilities

No one can predict the future with precision. Instead, we think in terms of scenarios. How the Grey Rhinos play out will determine which of these scenarios becomes reality.

Base Case Scenario (60% probability)

In our central scenario, the transition from Phase 1 AI enablers to Phase 4 AI adopters continues, but with some friction.

Energy constraints, grid bottlenecks, labor shortages, and organizational resistance all slow things down. Some projects are delayed. Some budgets are revised. But overall, progress continues. The leading companies in the \$GSXUPROD basket use this period to move beyond experiments and prove that AI can consistently improve their margins.

As these improvements become visible in quarterly earnings – higher revenue per employee, lower personnel costs as a share of revenue, and expanding operating margins – we expect the market to gradually re-rate these companies over the next two to three years.

Alternative (Bear) Scenario (40% probability)

In the bear scenario, the Grey Rhinos prove more stubborn and more damaging.

Power shortages persist longer than expected. Transformer shortages and grid constraints cause repeated project delays. At the same time, a credit event – possibly involving a highly leveraged AI infrastructure player such as \$ORCL – shakes confidence in the entire AI CapEx cycle.

In this environment, boards and CFOs may decide to slow or cut back their most ambitious AI investment plans. Pilot purgatory, instead of being a temporary phase, becomes the default. Productivity gains arrive later, in smaller amounts, and the market's enthusiasm for AI adoption fades, at least for a time.

Key Triggers to Watch

To understand which scenario is unfolding, we monitor several key signals.

We pay close attention to credit spreads, particularly the CDS spreads of companies like \$ORCL that have taken on significant debt to fund AI projects. We also watch broad indicators such as high-yield bond ETFs (\$HYG) for signs of rising stress.

We track infrastructure timelines, especially announcements from hyperscalers and utilities about delays or cancellations in data center and power projects. Multi-quarter delays are an important warning sign.

And we watch margin trends in the companies that make up \$GSXUPROD. If, after several quarters of investment, we see little or no improvement in operating margins, that suggests the productivity thesis is not yet materializing as expected.

7.0 The MoatPeak Perspective: How We Interpret the Shift

At MoatPeak, we try to do more than simply describe what is happening. Our aim is to interpret events using clear mental models and second-order thinking, so we can understand how today's developments translate into tomorrow's opportunities.

We see Phase 4 AI as a change in the very way value is created and recognized in markets.

From Slides to Free Cash Flow

For several years, companies were rewarded just for having an "AI story." Adding a slide about AI to an investor deck, or dropping the word "AI" on an earnings call, was often enough to move a stock.

That period is ending. Investors increasingly want proof, not promises. They are asking: Does AI show up in free cash flow? Does it show up in margin expansion? Does it improve the stability and quality of earnings?

This shift in expectations naturally favors companies that are using AI to create immediate operational leverage – fewer people needed for the same output, or more output with the same number of people. It is less favorable for capital-intensive builders that are still burning cash on large-scale infrastructure, hoping for a payoff further down the line.

The Railroad Analogy

The history of railroads offers a helpful parallel. In the early days, huge fortunes were made by companies that built tracks, bridges, and tunnels. But over time, much of the enduring value shifted to the companies that operated trains, moved goods, and created entirely new business models around logistics and transport.

In AI, companies like NVIDIA (\$NVDA) and other infrastructure providers have played the role of track-layers in Phase 1. In Phase 4, we believe companies like Walmart (\$WMT) and JPMorgan (\$JPM) are positioned to become the new magnates. They are not just using AI to cut costs; they are using it to rethink how they manage inventory, move goods, interact with customers, and allocate capital. (*see Figure 13*)

The Final Insight: Building the Railroad vs. Building the Economy

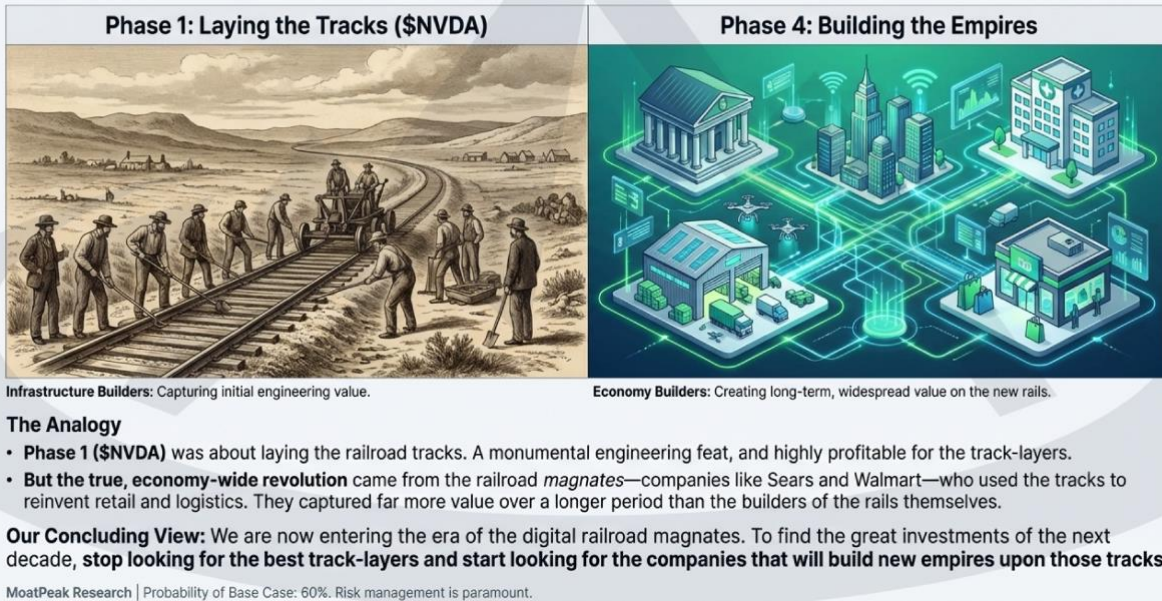


Figure 13. Railroad analogy: tracks versus the economy built on top

Key takeaway: Over time, value can migrate from infrastructure builders to the operators who compound returns by using the new 'tracks' effectively.

History suggests that, over the long run, more of the economic value tends to flow to the operators who build on top of the infrastructure than to the initial builders of that infrastructure.

AI as a General-Purpose Technology

Finally, it is important to recognize that AI is no longer confined to the “technology sector.” It is becoming a universal factor of production, like electricity, computing, or broadband connectivity.

Over the next decade, the real question will not be “Is this a tech company?” Instead, it will be, “How deeply is AI embedded in this company’s daily operations?” and “How strong is its moat built on data, algorithms, and organizational capabilities?”

The companies that win will be those that can consistently gather and use data, maintain and improve their models, and manage the human side of change – training people, redesigning processes, and updating incentives – so that AI becomes part of the firm’s DNA rather than a series of side projects.

8.0 An Actionable Framework for the Retail Investor

The analysis so far has focused on what is happening and why it matters. The natural next question is: what can an individual investor do with this information?

This is not personal investment advice, but it is a way to structure your own thinking. (see Figure 14)

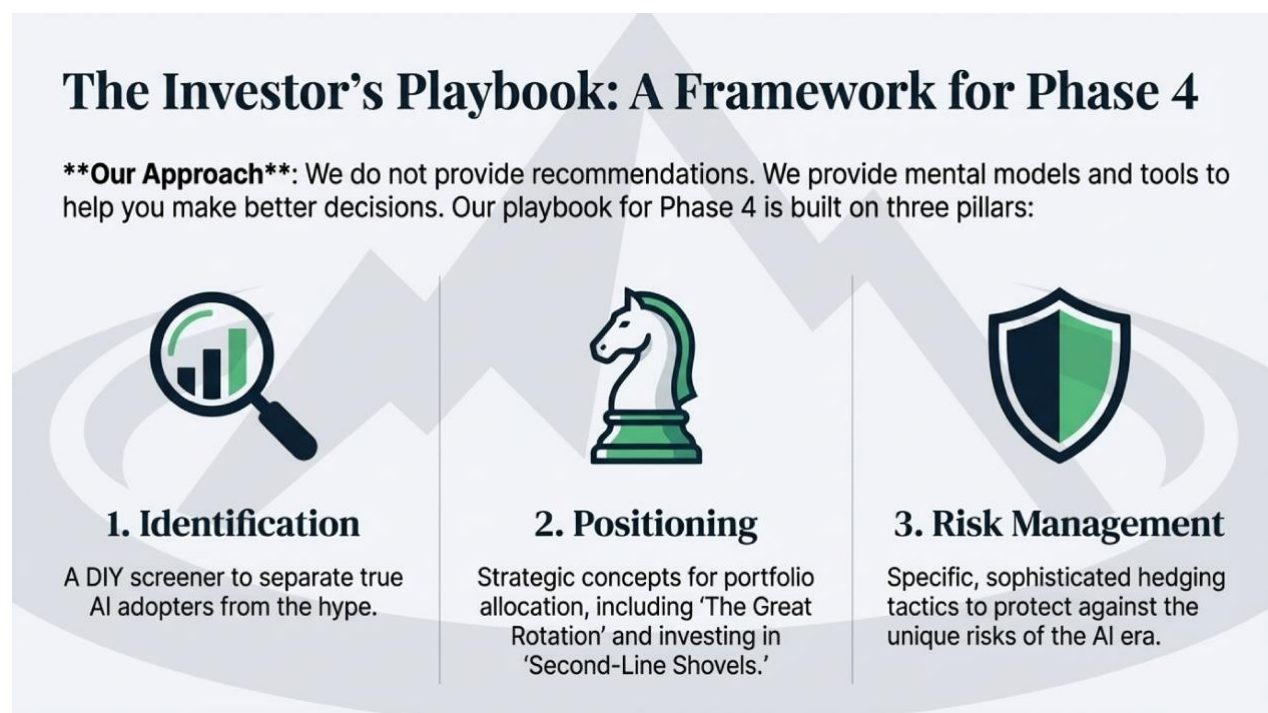


Figure 14. A Phase 4 playbook: identification, positioning, risk management

Key takeaway: Our framework emphasizes screening for true adopters, positioning for the rotation, and explicit risk management for shocks.

One helpful step is to develop your own simple “AI adopter” checklist. For any company you are considering as a Phase 4 beneficiary, look at how key metrics evolve over time. Are revenue, EBIT, free cash flow, and capital expenditure per employee improving? Is the share of revenue spent on personnel gradually edging down, or at least not rising as fast as revenue? Are gross margins stable or growing?

If AI is delivering real productivity gains, some of these improvements should eventually show up in the numbers. For software-focused businesses, you can also watch metrics such as net revenue retention (NRR) and average revenue per user (ARPU), which can indicate whether AI is helping deepen customer relationships and increase the value of each user.

A second element is to think about capital rotation. In broad terms, many Phase 1 beneficiaries, such as semiconductor names and related ETFs like \$NVDA and \$SMH, have already experienced very strong runs. At the same time, many Phase 3 and Phase 4 adopters, including software and traditional companies in ETFs like \$IGV, remain less richly valued relative to their potential AI upside.

Conceptually, an investor might choose, over time, to trim exposure to the most extended Phase 1 names and redeploy some of that capital into carefully chosen Phase 4 adopters with strong balance

sheets, solid free cash flow, and clear evidence that they are using AI to strengthen their competitive positions. We think of these as “quality value” names enhanced by technology. (see Figure 15)

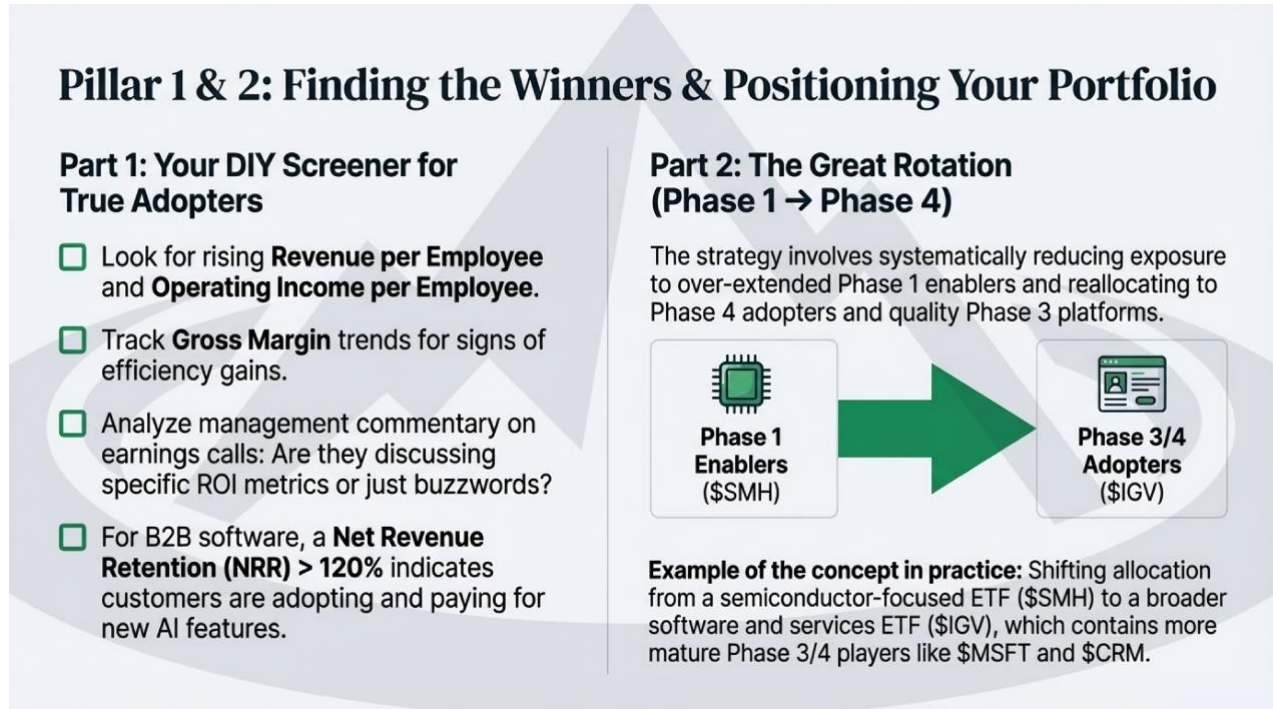


Figure 15. Finding winners and positioning: adopter screen and rotation concept

Key takeaway: A practical approach is to track productivity and margin metrics and consider rotating from extended Phase 1 exposure toward quality Phase 4 names.

A third consideration is risk management. Major technological shifts are rarely smooth. They are noisy, emotional, and sometimes messy. Periodic scares around AI regulation, misuse, security, or ethics are likely. Having some form of protection in place can make it easier to stay invested instead of being forced out by volatility at exactly the wrong time.

For some investors, that might mean using instruments tied to market volatility, like VIX options or ETFs such as \$VIXY, as a form of insurance against sharp spikes in fear. For others, it might mean maintaining a stable allocation to defensive sectors like healthcare, for example via \$XLV, which have historically been more resilient when growth or tech stocks come under pressure. Investors particularly concerned about credit risk could also consider tools such as put options on high-yield bond ETFs like \$HYG as a way to hedge against an AI-related credit shock.

All of this should be tailored to individual circumstances, of course, but the broad idea is simple: you do not need to choose between participating in the AI opportunity and managing risk. With a thoughtful approach, you can do both.

The question for 2026 is no longer whether AI will create value. That debate is over.

The real question is who will capture that value.

Will it be the companies that built the digital tracks, or the companies that learn to run the most efficient, resilient, and profitable operations on top of them? Our analysis suggests that the greatest fortunes of this AI era are still ahead of us, and that they will belong disproportionately to the high-quality adopters – the firms that can integrate AI deeply into their businesses and turn it into durable, compounding cash flows. *(see Figure 16)*

MoatPeak — Independent Research

MoatPeak is an independent macroeconomic and sector-focused research platform dedicated to helping investors navigate complex market regimes through scenario-based analysis, institutional synthesis, and disciplined risk frameworks. Our research integrates insights from global investment banks, asset managers, and macroeconomic data sources, translated into clear, actionable perspectives for long-term, strategic investors.

MoatPeak is not affiliated with any financial institution, broker-dealer, or asset manager.

Important Disclosures & Disclaimer

This report is provided for informational and educational purposes only and does not constitute investment advice, a recommendation, an offer, or a solicitation to buy or sell any securities, financial instruments, or investment products. All views expressed reflect the author's judgment at the time of publication and are subject to change without notice. Past performance is not indicative of future results. Investing involves risk, including the potential loss of principal. MoatPeak does not provide personalized investment advice and does not consider the investment objectives, financial situation, or specific needs of any individual reader. Readers are encouraged to conduct their own research and consult with professional advisors before making any investment decisions.

While the information contained herein is believed to be reliable, it is derived from publicly available sources and third-party materials. MoatPeak makes no representation or warranty as to the accuracy, completeness, or timeliness of the information provided.

Intellectual Property & Usage

© 2025 MoatPeak. All rights reserved.

This report, including its content, structure, analysis, charts, and visual elements, is the intellectual property of MoatPeak and may not be reproduced, redistributed, or transmitted, in whole or in part, without prior written permission.

Contact

research@moatpeak.com www.moatpeak.com



Figure 16. Important disclosures and disclaimer

Key takeaway: This report is for informational purposes only; readers should conduct their own research and consider professional advice.