FIERA CAPITAL

JULY 2026

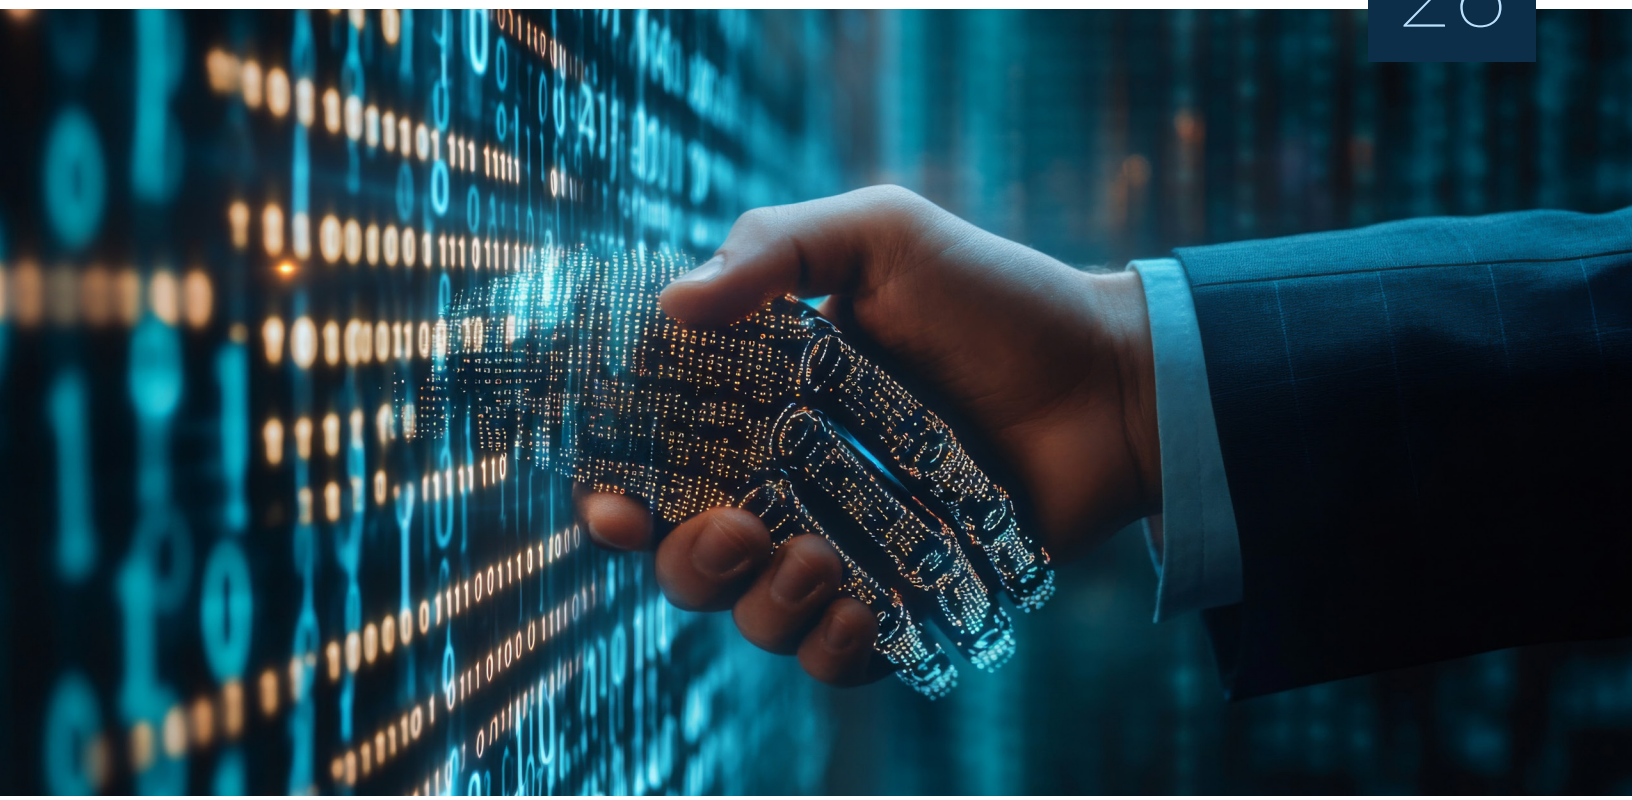

The AI Platform Shift

A framework for navigating evolving moats and business models in the AI era

Andy Gardner
Senior Portfolio Manager,
Fiera Atlas Global Companies Team

Key Takeaways

- **AI is a general-purpose platform shift:** adoption follows a sequence (infrastructure first, workflow deployment second, broader organisational redesign later) with value now migrating toward the physical and operational bottlenecks that enable scalable, cost-efficient deployment.
- **Moats are being redefined:** parts of the stack grounded in physics and process complexity may have stronger positioning despite greater cyclicality; software is not disappearing, but will be rebuilt as organisations redesign work around AI, with agents accelerating that shift.
- **Discipline matters more as the end state is uncertain:** we have reallocated capital toward areas of strengthening competitive advantage and away from harder to underwrite exposures. Refresh has delivered a portfolio that is demonstrably higher in quality, broader in diversification and governed by tighter process control.

Introduction

Over the past decade and a half since the global financial crisis, technology stocks have been an engine of growth in world equity markets. From 2015 to 2025, the MSCI All Country World Information Technology Index (MSCI ACWI IT) delivered about 25% annualized returns, far outpacing the roughly 13% annualized return of the MSCI All Country World Index (MSCI ACWI).

That outperformance, however, masks an equally important truth – periods of exceptional innovation frequently trigger sharp drawdowns and rapid sentiment shifts. The tech sector suffered an 80% collapse in the early 2000s after the dot-com boom, compared to a 55% decline for global equities, and has experienced twice as many 10% drawdowns as the broader market over the past two decades.

This is because technological progress, while a driver of long-term growth, is also inherently disruptive, leading to clear winners and losers across each generation of technology. From mainframes to personal computers, from PCs to the internet, from the web to smartphones and from on-premise software to the cloud, every platform transition has created enduring winners and notable casualties.

Risk-Return Profile of Global Technology Equities

~25%
MSCI ACWI IT 10 year annualized return vs ~13% for MSCI ACWI overall

16
Number of drawdowns greater than 10% in MSCI ACWI IT over the past 20 years vs 7 drawdowns for MSCI ACWI overall

-80%
Peak to trough decline during 2000–2002 (30 months) vs –55% for MSCI ACWI

Source: Bloomberg; As of June 30, 2026

Structural Shifts in Technology Platforms

Source: Adapted from Ben Evans https://www.ben-evans.com/presentations

AI is the next major step in technology's evolution. But it is distinct in two respects. First, it is a true general-purpose technology, meaning its impact is likely to be pervasive across industries, more akin to electricity than to a conventional product cycle. Second, its eventual end state remains unclear, creating an unusually wide range of possible outcomes. That increases the need for disciplined stock selection. The right response to this changing landscape, in our view, is to do what technology itself does – evolve.

As quality-oriented investors, the philosophical principles around owning exceptional companies remain sacrosanct. But each platform shift reshapes which moats strengthen, which weaken and which business models prove most resilient. The investable universe adapts. Our review of the evolving technology landscape has therefore been substantial. Over several months, we have mapped these moving parts into an investment framework designed to capture technology's growth while remaining selective as business models and competitive positions evolve. This paper is a condensed summary of that work.

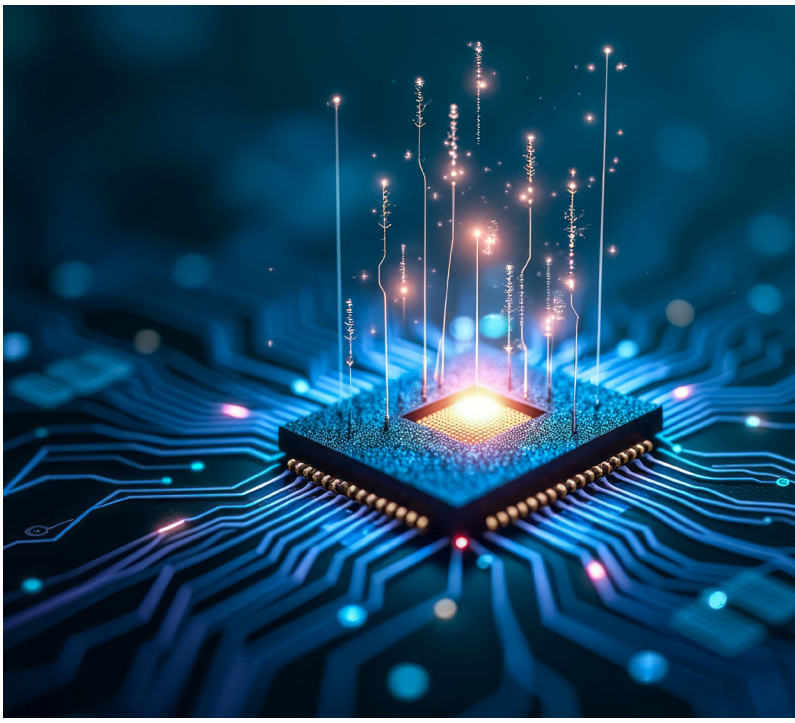

AI is a Sequence: Infrastructure First, Workflow Reorganization Second

The last technology era (roughly 2010 to 2022) was defined by cloud computing and software-as-a-service (SaaS) business models.

Low inflation and declining interest rates boosted the present value of long-duration revenue streams, favoring companies with subscription-based, capital-light approaches. Software and internet businesses delivered recurring revenue with high incremental margins; they were rewarded with premium valuations and rising returns on capital. In contrast, hardware and semiconductor firms, with their capital intensity and more cyclical demand patterns, were often discounted by investors despite their foundational importance.

The rise of AI is now challenging these assumptions and unfolding in distinct phases. From an investment perspective, the infrastructure build-out is occurring today and remains the most visible phase of the AI transition. The more difficult question is how quickly organisations can redesign workflows, decision-making processes and operating models to fully utilise these new capabilities. While infrastructure adoption can be measured through capacity additions and capital expenditure, workflow transformation depends on organisational change, governance, incentives and employee adoption. As a result, the productivity benefits of AI are likely to emerge more gradually and unevenly than the underlying infrastructure investment cycle.

Recent evidence suggests that enterprises are progressing beyond experimentation and pilot programs toward real deployment in functions such as software development, customer service, knowledge management, cybersecurity and internal operations. However, the primary bottleneck is increasingly organisational rather than technical. Data access, governance, model selection, change management and workflow redesign are becoming the critical constraints on adoption. This distinction matters because technological capability can improve much faster than organisations can redesign the way work is performed.

Layered View of the Technology Stack

Layer 1
Semiconductors, Equipment and Manufacturing

AI and the Limits of Physics – AI is turning advanced compute hardware into a strategic resource. Performance gains increasingly depend not just on chip design, but on process technology, specialized equipment, advanced packaging, memory hierarchy and thermal management. Key constraints include maintaining yields at leading-edge nodes, expanding capacity fast enough to meet demand and managing power density and cooling in larger data centres. Moats are deepest where companies combine extreme R&D intensity, physics-constrained engineering know-how, long qualification cycles and large installed bases that create ecosystem lock-in. These factors entrench a small number of systems-relevant leaders across the semiconductor supply chain.

Layer 2
Networking and Interconnect

Networks as Compute Multipliers – In large-scale AI clusters, networking has become as important as the chips themselves. As deployments scale, bandwidth, latency and reliability increasingly determine system performance and cost. Key challenges include delivering very high throughput with low latency, managing power consumption and maintaining reliability as systems grow in size and complexity. The strongest positions sit with vendors that combine high-performance silicon, advanced optics and system software into integrated solutions. Their moats are reinforced by deep customer relationships, technical breadth and a proven ability to operate at hyperscale.

Layer 3
Cloud 2.0 (AI-Optimized Cloud Infrastructure)

From General-Purpose Cloud to AI-Optimized Systems – The 'Cloud 2.0' era is defined by intelligent, AI-tuned infrastructure rather than generic computing services. The leading cloud providers are developing vertically integrated stacks, from custom AI chips and high-speed networks to specialized storage and data services, aimed at maximizing performance per dollar for AI workloads. Differentiation in cloud is shifting toward system-level optimization, advanced orchestration software and the ability to help enterprise customers integrate and govern AI at scale, rather than simply offering raw computing power for rent. Key constraints at this layer include the need to build out massive data centre power and cooling capacity, maintain secure supply of specialized hardware (especially advanced processors and accelerators) and orchestrate complex AI models and data flows reliably for large clients. Hyperscalers with global infrastructure, strong developer ecosystems and deep enterprise ties are poised to excel in this environment – they can deliver cost-effective AI performance, broad solution portfolios and the support businesses require to implement AI-driven transformations.

Layer 4
Data, Security and Platforms of Trust

Governance and Protection in an AI-Driven Enterprise – As AI moves into production, data governance and security become even more important. Enterprises must manage data quality and lineage, meet privacy and compliance requirements, control identity and access and detect new threats created by AI adoption. Platforms that provide identity, telemetry, policy enforcement and incident response can build strong moats because they occupy critical control points in enterprise IT. By turning data and policy into better security and reliability outcomes, they become increasingly indispensable as AI expands the reach and risk of digital systems.

Layer 5
Application Software and Workflow Automation

AI in the Flow of Work – As AI agents become embedded in enterprise workflows, software consumption is changing. Instead of users navigating applications through a graphical interface, AI may increasingly execute tasks via APIs and behind-the-scenes orchestration. This shifts value toward workflow ownership, permissioned data, deep integration and outcome-based pricing rather than seat count alone. The strongest software businesses will remain central to mission-critical processes and use AI to improve productivity and decision-making. Over time, they may evolve from systems of record into systems of context and action, retaining their importance by providing the secure environment, data and orchestration layer for AI to operate. By contrast, software with shallow integration or undifferentiated features will find it harder to defend its value.

Layer 6
Edge and Industrial Integration

AI in the Physical World – The next frontier for AI is at the edge and in industrial settings, where digital intelligence is applied to physical operations such as factories, supply chains, vehicles and connected devices. These use cases require low latency, high reliability and strong safety controls. Key challenges include delivering meaningful compute on constrained hardware, meeting regulatory standards and integrating complex hardware and software reliably in the field. Companies that combine proprietary hardware, specialised software and field deployment capabilities have an advantage. By controlling both the device and operating layer and by collecting unique real-world data, they can build self-reinforcing moats in their domains.

Not every company benefiting from AI will be a good investment. Growth alone does not ensure value creation. Many 'picks and shovels' businesses will attract capital but may struggle to earn strong returns if competition rises or supply outpaces demand. Our focus remains on companies with durable competitive advantages that can sustain superior economics through cycles and technological change. In practice, that means avoiding areas where industry structure drives price-taking and oversupply, such as commodity memory, undifferentiated compute rental or offerings without customer lock-in or network effects. The opportunity set may broaden in AI, but we believe only a subset of companies will meet the bar for durable moats and attractive long-term returns.

Software and the Reorganization of Work

The software debate is often framed too narrowly around whether AI will disrupt software. In our view, that is the outcome, not the starting point. The more important question is how organizations reorganize work around AI – how tasks are broken down, where permissions and control sit, how outputs are verified and which platforms remain central to execution.

The distinction is important because software is increasingly evolving from a product designed primarily for human users into a system that must also support autonomous and semi-autonomous agents. As AI capabilities improve, the interface itself may become less important than the workflows, permissions, governance structures and data assets that sit underneath it. Traditional SaaS bundled data, business logic, analytics, workflow, permissions and a graphical interface into a pre-built application that users navigated directly. In an AI-driven world, that model starts to change.

A simple way to think about it is this – an agent does not need a dashboard. A human user may log into five different applications and navigate dozens of screens. An agent can potentially complete those tasks through APIs, tools and governed-access layers without relying on the interface at all. In that world, consumption shifts from seats and UI time toward API calls, tasks completed, outcomes delivered and governed delegation.

But the bigger point is that this is not only a software redesign story. It is a firm redesign story. As organizations rethink how work is divided between people and agents, how decisions are reviewed and where governance, permissions and accountability sit, software durability will increasingly depend on whether a product remains central to that redesigned operating model.

This is why even strong software franchises may need to rebuild parts of their operating model. The issue is not that all software becomes obsolete, but that product architecture, distribution, pricing and cost bases may all need to adapt. That should not be surprising. Software has already gone through multiple reinventions – from mainframe to PC, from PC to internet and from on-premise to cloud. AI may prove to be the next such transition.

There is also a budget implication. AI spend is not arriving in a vacuum. In many cases, it is competing directly with existing (and large) software budgets, while vendor consolidation is becoming a more active priority for CIOs. That makes contract renewal, pricing power and product relevance more important than they have been for years.

Our conclusion is not that software disappears, but that software will be reshaped by how firms reorganize work around AI, much as it has been transformed in prior platform shifts.

We believe a more important question is who owns the learning loop? The model provider or the organisation? As humans and AI systems increasingly work together, organisations generate valuable proprietary assets in the form of prompts, evaluations, corrections, workflow traces and institutional knowledge. Over time, competitive advantage may depend less on access to a particular model and more on ownership of these learning loops, which allow organisations to continuously improve how AI is applied within their own environments.

That is why we are seeing a growing number of enterprises adopting multi-model architectures rather than standardising on a single AI provider. Frontier models may be used for complex reasoning tasks, smaller models for routine workflows, specialised coding models for software development and private models for sensitive internal use cases. In this environment, the orchestration layer that governs model selection, permissions, workflow execution and cost management may become more strategically important than any individual model provider or software provider.

Technology Cycles

When investing in technology, it is essential to consider the interaction of four overlapping cycles. These cycles help explain why outcomes can diverge widely in tech markets and underscore the need for careful portfolio management to support bottom-up stock analysis:

1 Structural cycle (multi-decade platform shifts)

This is the long-wave adoption and a source of great change – a new platform becomes viable, infrastructure is built and usage compounds along an S-curve. Importantly, value often migrates to different layers of the stack over time. Early in a shift, scarce inputs and enabling layers can capture outsized value; later, once capacity is abundant and cheap, value accrues to layers that control distribution, context and user workflows. During the late 1990s telecom boom, networks laid enormous fibre capacity ahead of demand and large portions remained unlit ('dark') for years. TeleGeography estimated in late 2002 that only around 10% of long-haul fibres in Europe and North America carried signals, highlighting the scale of the overshoot. That overbuild eventually became the backbone for the modern internet, a reminder that long-run utility can be real even when near-term economics are poor.

2 Industry cycle (supply and demand dynamics, ~2–5 years)

Many tech subsectors experience traditional cyclical swings in supply, demand and pricing. Surges in capital expenditure and production often lead to oversupply, inventory buildups or sudden falls in pricing and orders, even when underlying demand remains on a long-term growth path. Semiconductors, memory, networking hardware and other infrastructure often follow these capex cycles. A powerful secular trend like AI does not immunize companies from these shorter-term cycles if capacity is added faster than demand or if customer budgets pause. The key risk is not that demand disappears, but that periods of intense competition or oversupply can drive down returns before market balance is restored. This is why 'theme exposure' without a defensible moat can be value-destructive even in fast-growing markets.

3 Valuation cycle (market expectations and interest rates, ~1–3 years)

In technology, valuation often moves ahead of fundamentals. Multiples expand when investors extrapolate the structural narrative, when near-term momentum is strong or when discount rates fall. They contract when expectations reset, when rates rise, or when evidence emerges that growth is more cyclical, more competitive, or more capital-intensive than assumed. The most damaging drawdowns typically occur when a stock is priced for perfection and the industry cycle turns, that is, when investors pay a peak multiple on peak earnings. This 'double reset' was a defining feature of many high-duration growth stocks as conditions tightened after the 2021 peak.

4 Market cycle (short-term positioning and flows, weeks to months)

Transient market dynamics (momentum trading, algorithmic strategies, crowded positioning and rapid shifts in risk appetite) can move technology stocks around in the short run. Quarterly earnings surprises or shifting narratives (like the latest AI trend) can trigger abrupt inflows or outflows, sometimes pushing stock prices far beyond or below what long-term fundamentals would imply. This volatility is a double-edged sword – it can render strong businesses temporarily overvalued and vulnerable to correction but it can also create buying opportunities when sentiment swings too negative relative to a company's underlying fundamentals.

Where does AI stand across these cycles?

Broadly, AI remains in the early stages of its structural adoption cycle, although adoption has accelerated significantly over the past year. The debate is increasingly shifting from whether enterprises will use AI to how they can deploy it safely, economically and at scale. Adoption is now evident across software development, customer support, document processing, cybersecurity, knowledge management and research workflows. However, broad organisational transformation remains at an early stage. The primary barriers are no longer model capability alone. Governance, workflow integration, data access, security, organisational incentives and change management are increasingly determining the pace of adoption.

This distinction is important for investors. The gap between theoretical capability and realised business value remains substantial. While AI systems are advancing at an extraordinary pace, organisations typically adopt new technologies far more gradually. As a result, the long-term opportunity remains significant, but the path from technological capability to economic value is likely to be uneven across industries and companies.

The AI Adoption Gap Across the Labour Market

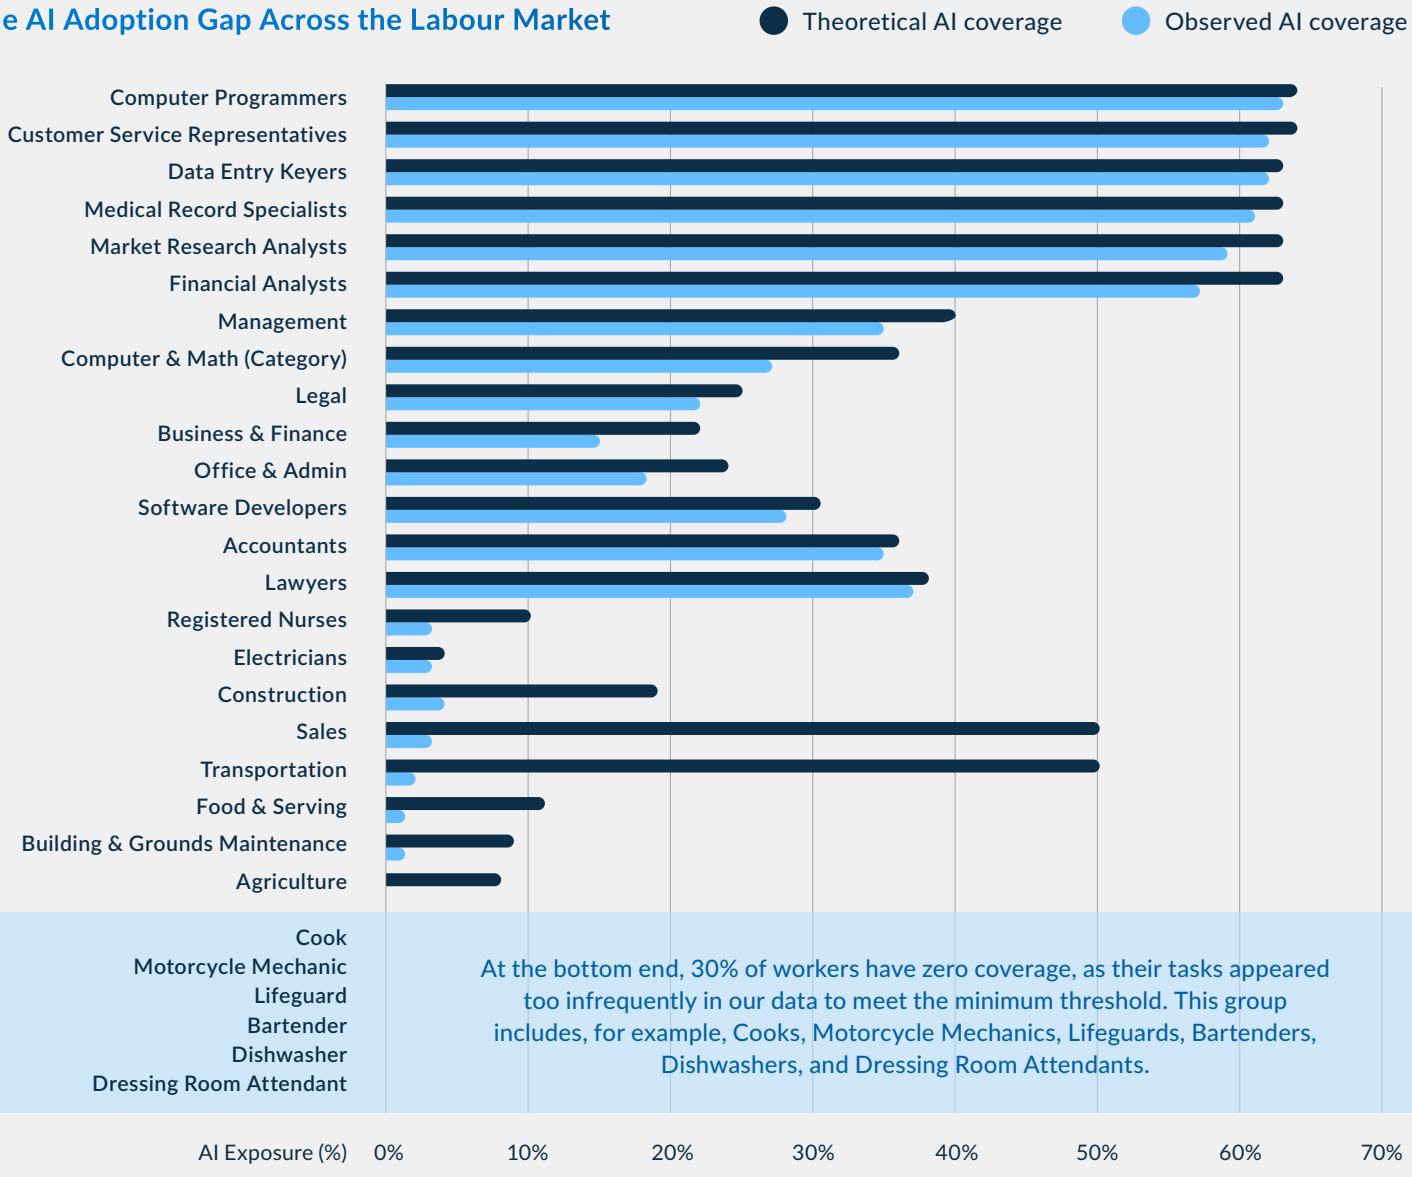

Source: Anthropic https://www.anthropic.com/research/labor-market-impacts; As of March 5, 2026

Across the industry cycle, parts of the AI 'enabling' stack are well into a cycle upswing – hyperscaler cloud providers, semiconductor companies and equipment makers are investing heavily to relieve short-term bottlenecks in compute, networking and manufacturing capacity. However, even at this foundational phase – where there is great concern a repeat of the dark-fibre days of the dot-com bubble – the build-out does not yet seem excessive. GPUs are not sitting 'dark' like fibre was in 2002. Compute is being used with utilization close to 100%, and demand is outstripping supply, which is why prices are rising. The more important question is not whether capital intensity remains high, but whether the leading platforms can earn attractive returns on that larger capital base.

A key part of our AI work has been to analyze return on invested capital (ROIC) scenarios in depth. The strongest hyperscalers are not merely buying accelerators and reselling compute. They are using proprietary silicon, software control planes, distribution and enterprise workflow integration to lower cost per unit of intelligence on the one hand, while monetizing that stack at a much higher level of value-added on the other. While we must monitor capital discipline closely, the way in which they are managing this 'value spread' so far may suggest the multi-decade trend of technology companies benefitting from increasing returns to scale may still be intact.

On the valuation front, many leading technology stocks have seen their valuations moderate from 2021 highs, narrowing the premium that 'big tech' trades at relative to the broader market. And from a market sentiment standpoint, frequent rotations and volatility around AI news and earnings announcements suggest that short-term speculation remains elevated.

Facing these overlapping cycles, our strategy aims to avoid stacking risks – for example, a stock with a fragile moat, late-cycle exposure and a rich valuation. We want to invest only where moats are durable and where the combination of cycle, position and valuation is reasonable.

Global Technology Sector ROIC Over Time (1980–2025)

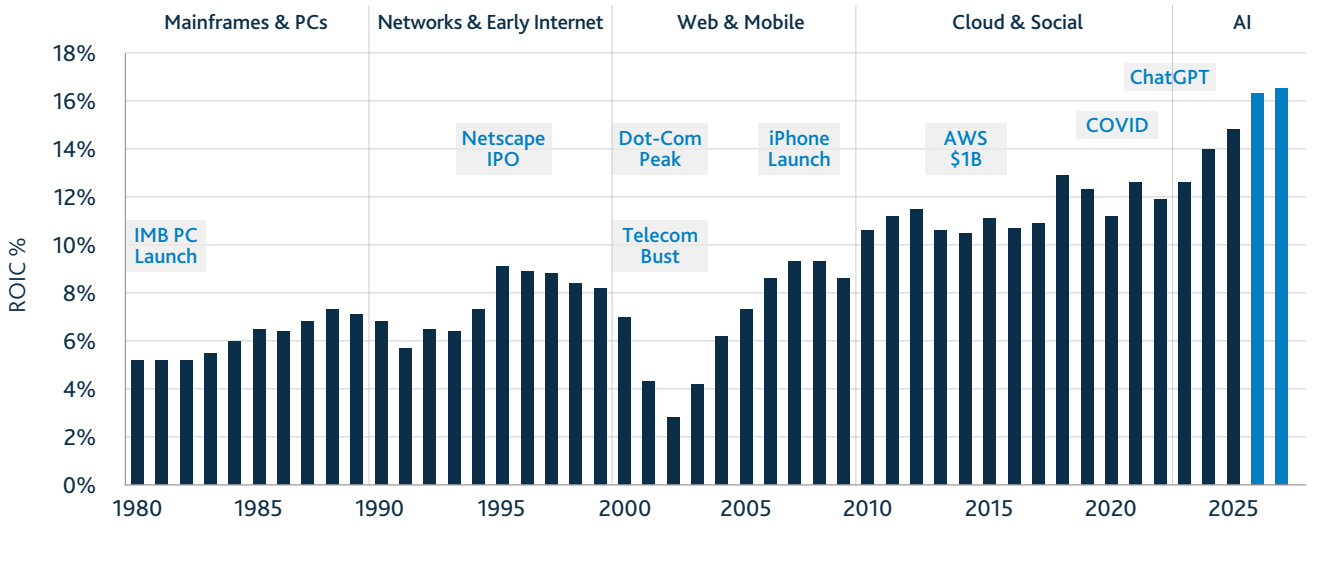

Source: Estimated from UBS HOLT chart (1980-1999); Bloomberg BDH for ROIC (2000-2025); As of December 31, 2025

Process and Portfolio Implications

Our exploration of the changing technology landscape has also coincided with a meaningful set of enhancements to our investment process. The issue is not simply that AI creates new opportunities; it also raises the importance of being explicit about which risks we are willing to underwrite.

Cyclical and valuation risks are usually observable and, to some extent, measurable. They can often be managed through position sizing, diversification and patience around mean reversion. Disruption risk is different. It is less a question of short-term volatility and more a question of thesis and terminal value – whether a company's moat, growth runway or relevance within the value chain may erode over time. That makes it harder to manage through traditional portfolio construction tools alone and necessitates a stricter test of investability.

We introduced a formal Monitoring Score for each holding to flag early signs of improving or deteriorating fundamentals and competitive position, so we can identify change and adjust position sizing more proactively. We also created a dedicated industry oversight function to maintain an up-to-date view of major technological trends and to track how our portfolio exposures align with broader market dynamics. Our framework itself is treated as a living set of hypotheses that we continuously test and refine. To support that work, and aligned to how AI is evolving the nature of work, we use specialized agents to challenge our assumptions, surface conflicting evidence and identify emerging changes, while always applying our own judgement.

We significantly rebalanced our technology exposure in early 2026. We reduced several positions that, while still high-quality businesses in absolute terms, faced a higher burden of proof because of greater disruption risk. We redeployed that capital into investments more closely aligned with the resilient layers of the stack identified earlier.

Following these moves, technology accounts for roughly one-third of our portfolio by net asset value. Equally important, we applied the same review to the remaining two-thirds of the portfolio beyond technology. The largest capital increases outside the sector have been in select industrials and specialist services businesses, where we see durable competitive advantages, strong structural demand and attractive prospective returns. These exposures complement our existing holdings in consumer, financial and healthcare franchises, resulting in a diversified portfolio of high-quality companies across industries, end markets and geographies. Many of these businesses are either relatively insulated from direct AI disruption or are well placed to benefit from adopting AI themselves.

Conclusion

AI is a transformative platform transition that will reshape business and society over time, but its full impact will unfold in stages rather than as a single event. Already, the shift toward an AI-centric world is altering the economics of technology, emphasizing the integration of hardware, data and intelligence, while forcing companies to rethink workflows and competitive strategy.

Our investment approach aims to harness the upside of this transition while guarding against its uncertainties. We are focusing capital on areas where competitive advantages are becoming stronger and where we see clear pathways to durable value creation, while maintaining caution in areas where disruption risk is elevated. The process changes we have implemented are all designed to help us navigate the coming volatility with discipline and agility.

fiera.com

Important Information

Fiera Capital Corporation ("Fiera Capital") is a global, independent asset management firm listed on the Toronto Stock Exchange (ticker FSZ) offering customized multi-asset solutions across public and private markets. In the United States, investment advisory services are provided only through Fiera Capital affiliates that are registered with the U.S. Securities and Exchange Commission ("SEC") or operate under an applicable exemption. Registration with the SEC does not imply a certain level of skill or training.

This document is confidential and intended solely for the recipient. It may not be shared, reproduced, or distributed without permission. Nothing in this document should be viewed as a recommendation, offer or solicitation to buy or sell any security or investment and does not constitute investment, legal, tax or accounting advice. Services are offered only to qualified investors and only in accordance with applicable laws and regulations in each relevant jurisdiction. Information is believed to be accurate as of the publication date but may change without notice; no warranty is provided, and Fiera Capital and its affiliates disclaim liability for its use.

Past performance is not a guarantee of future results. Inherent in any investment is the potential for loss. Target returns are aspirational, forward-looking and do not represent actual performance. There is no guarantee that such performance will be achieved, and actual results may vary substantially. Metrics (e.g., exposures, ratios, characteristics) are for reference only and may not capture all relevant factors. Different metrics may lead to materially different conclusions; Any specific holdings mentioned are for illustration only, and may not represent the full portfolio, past holdings, or future positions. Indices are unmanaged, not investable, and do not reflect fees or expenses. Index comparisons are provided for context and portfolio holdings may differ significantly from those of any index; All investment strategies involve risks, including market, economic, financial, operational, liquidity, valuation, and regulatory risks.

Certain strategies may use leverage, derivatives, or concentrated exposures, which can increase volatility and risk of loss; No strategy, diversification approach, or risk management technique can eliminate risk, or guarantee returns in all market environments; Investors should review relevant governing documents and consult their own advisers before making investment decisions.

Environmental, social, governance ("ESG") or impact related goals, commitments, or initiatives referenced in this document are voluntary, may not apply uniformly across strategies, and may be modified or discontinued at Fiera Capital's discretion. ESG-related processes do not guarantee any specific investment outcome.

THIS DOCUMENT IS ISSUED BY: Abu Dhabi Global Markets: Fiera Capital (UK) Limited which is regulated by the Financial Services Regulatory Authority; **Cayman Islands:** Fiera Capital (Asia), L.P. which is regulated by the Cayman Islands Monetary Authority; **European Economic Area ("EEA"):** Fiera Capital (Germany) GmbH which is authorized and regulated by the Bundesanstalt für Finanzdienstleistungsaufsicht; **Hong Kong:** Fiera Capital (Asia) Hong Kong Limited which is regulated by the Securities & Futures Commission of Hong Kong; **Singapore:** Fiera Capital (Asia) Singapore Pte. Ltd. which is regulated by the Monetary Authority of Singapore; **United Kingdom:** Fiera Capital (UK) Limited and Fiera Real Estate UK, which are authorized and regulated by the Financial Conduct Authority; **United States:** Fiera Capital Inc. Fiera Capital (UK) Limited and Fiera Comox Partners Inc. are registered as investment advisers with the Securities and Exchange Commission ("SEC").

Additional registration and licensing information available here: https://www.fieracapital.com/en/registrations-and-authorities.

Version STRENG005