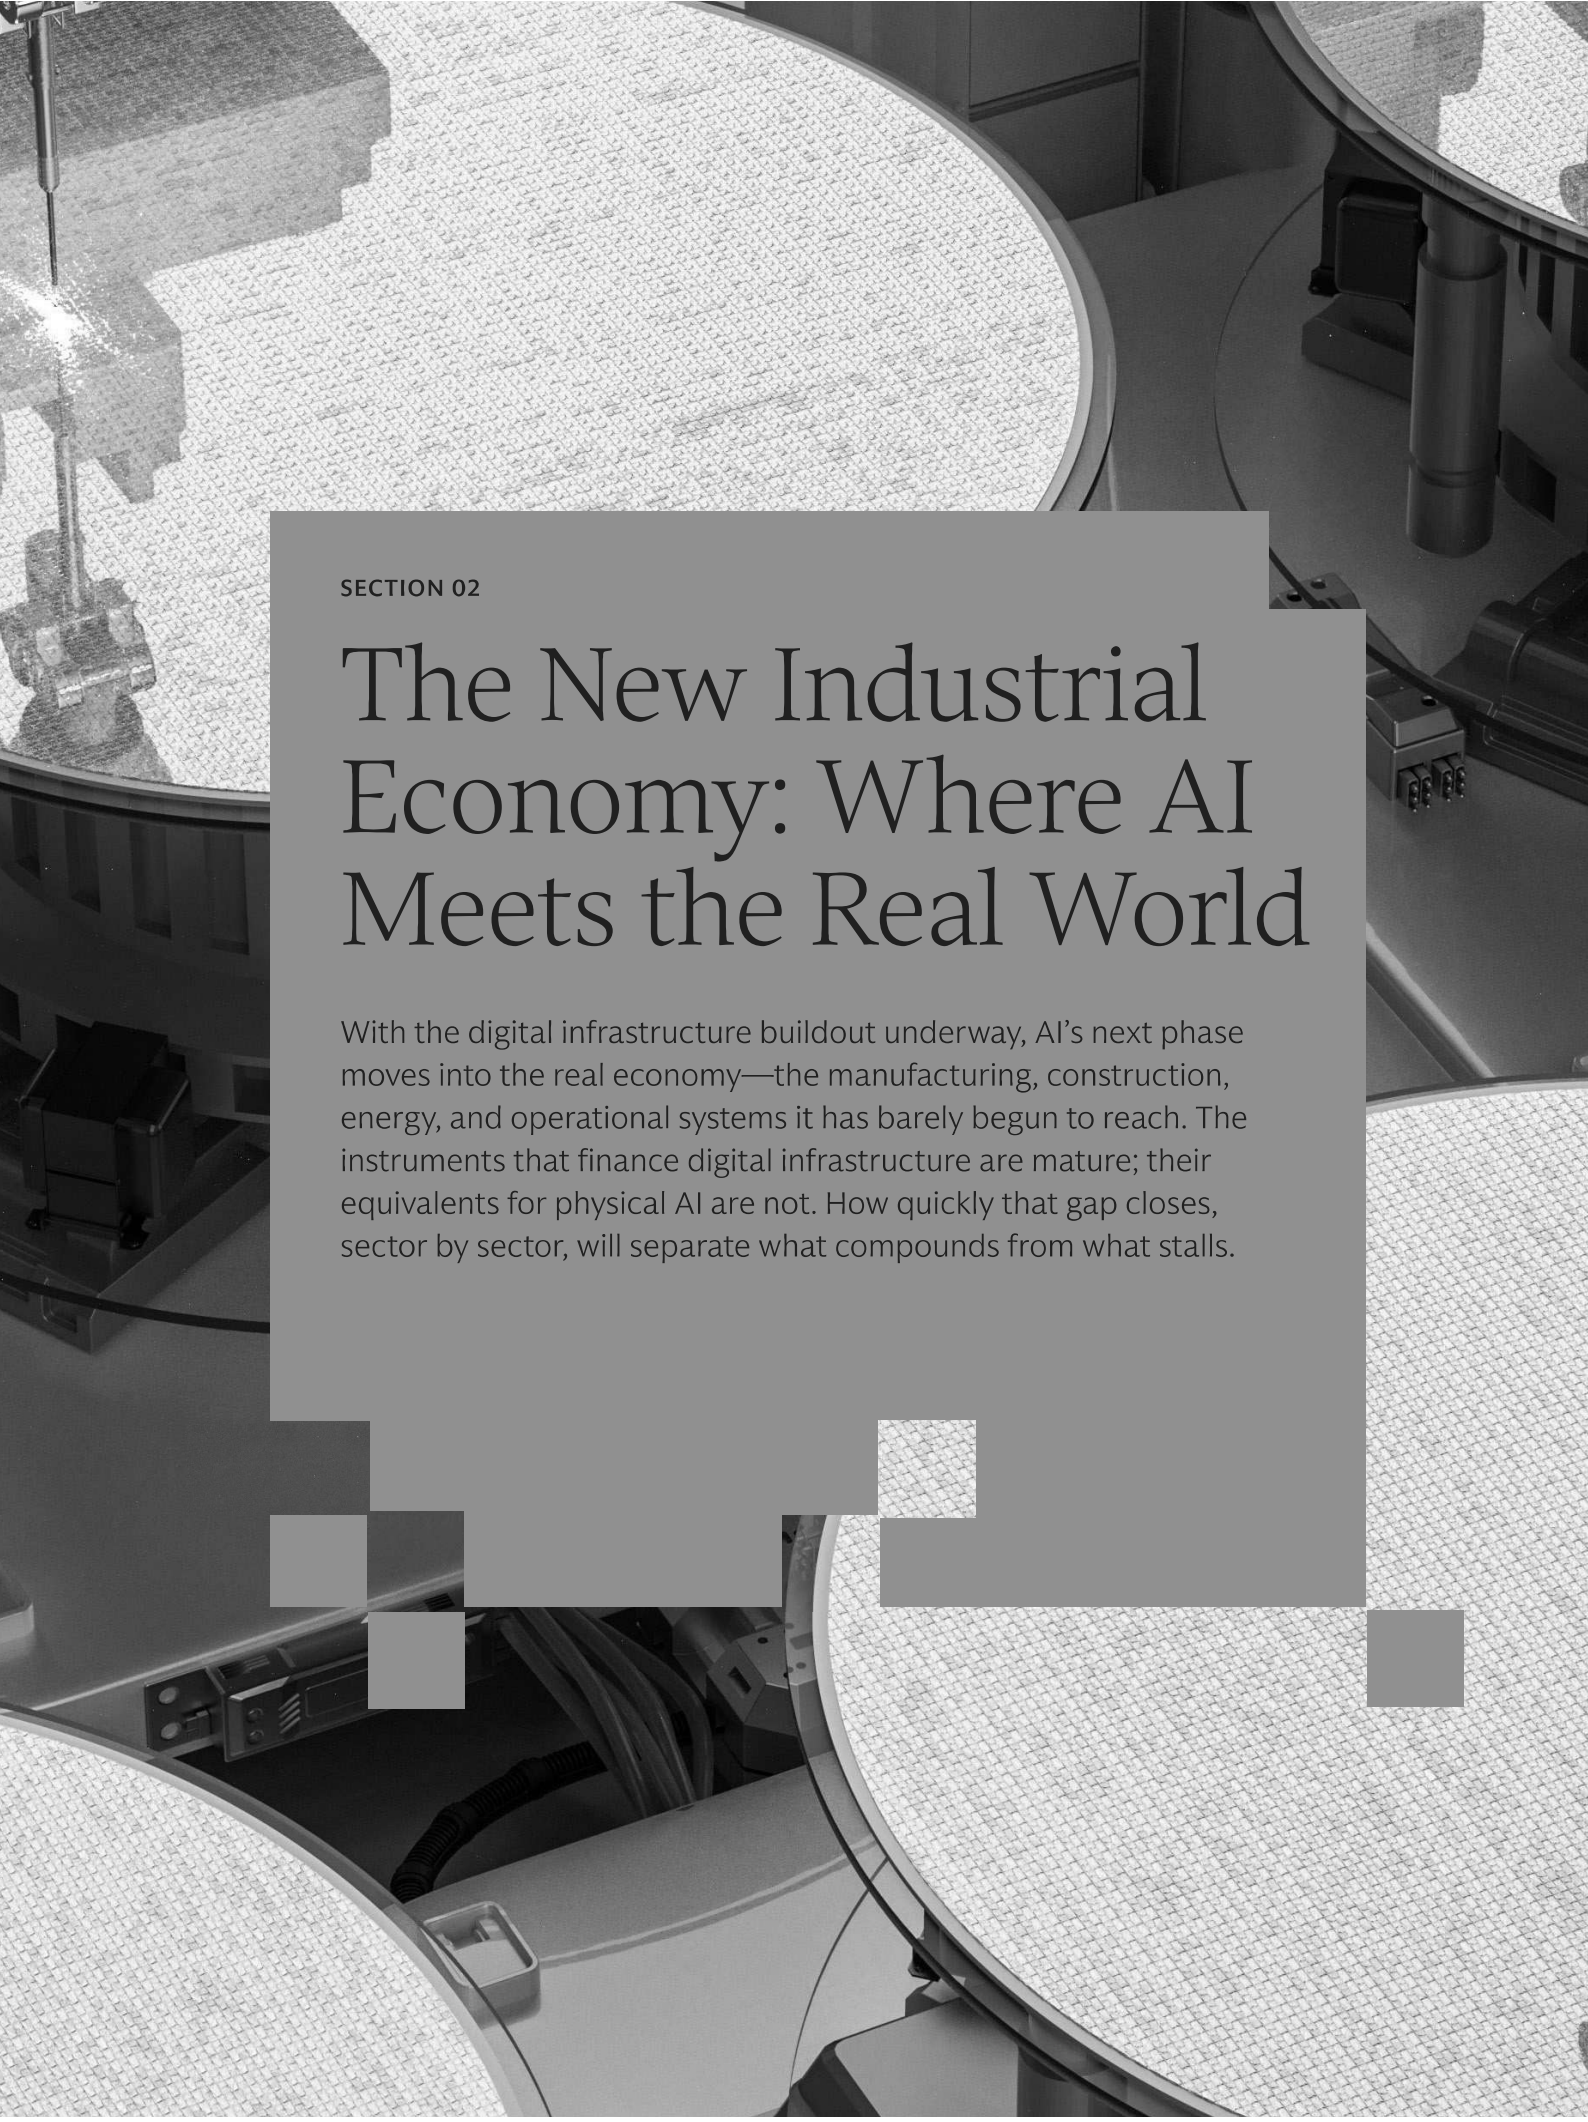


SECTION 02

# The New Industrial Economy: Where AI Meets the Real World

With the digital infrastructure buildout underway, AI's next phase moves into the real economy—the manufacturing, construction, energy, and operational systems it has barely begun to reach. The instruments that finance digital infrastructure are mature; their equivalents for physical AI are not. How quickly that gap closes, sector by sector, will separate what compounds from what stalls.

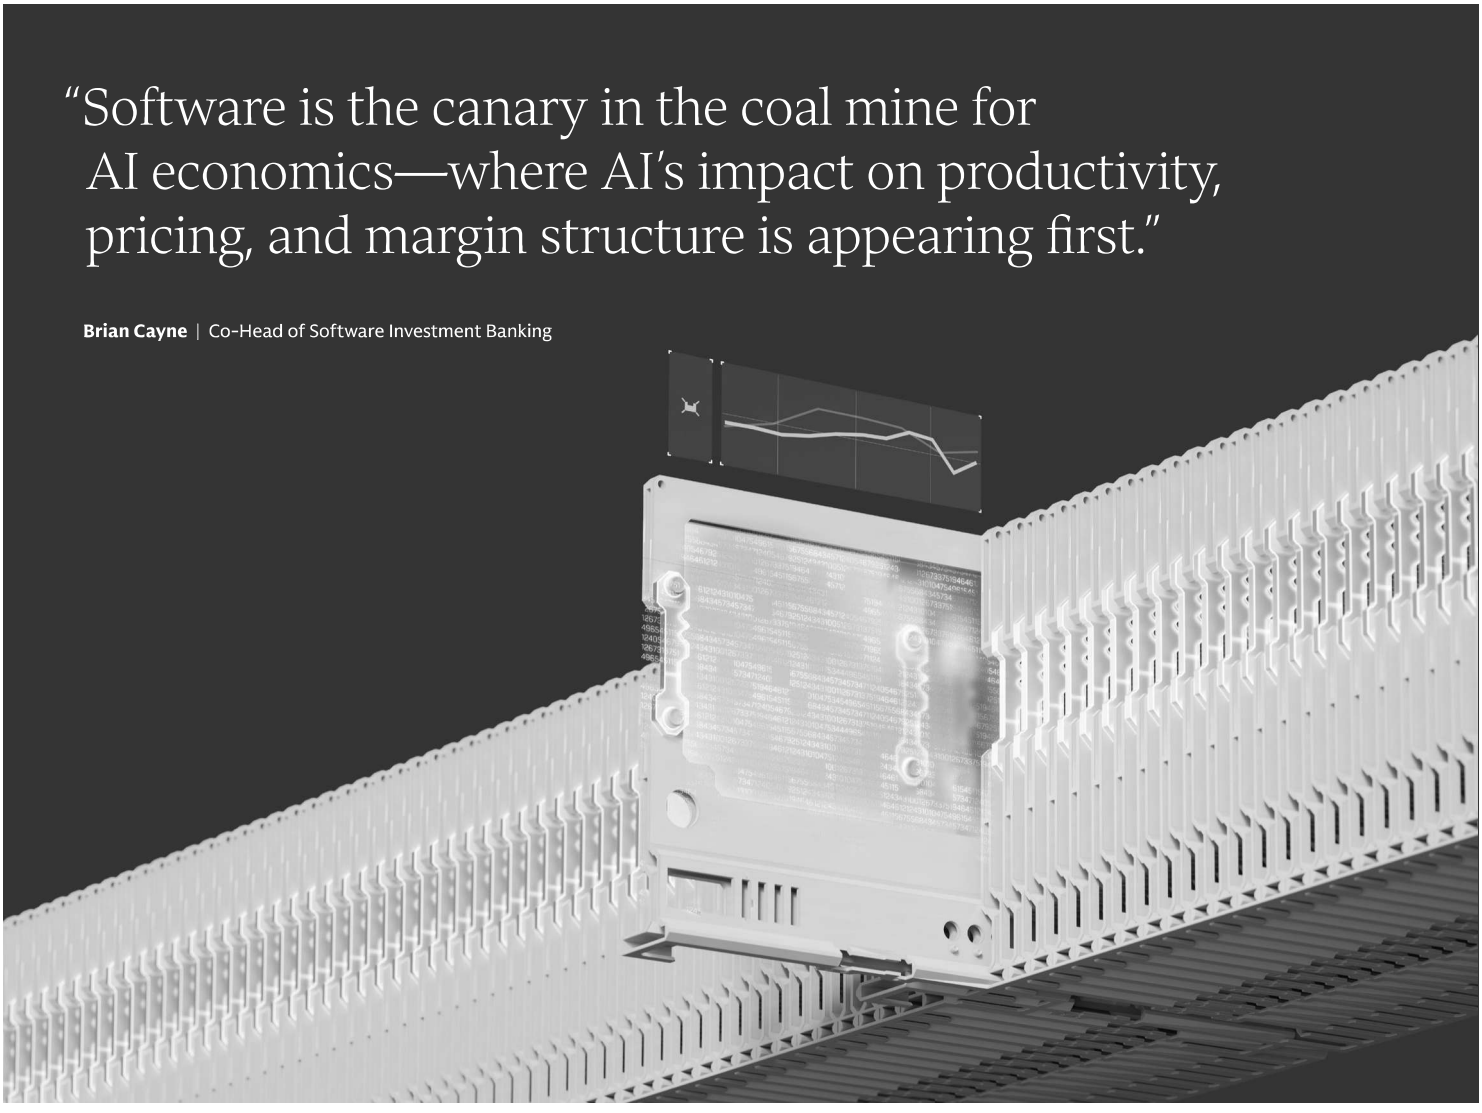
# I. Software: The First Domain and Enterprise Shock

Goldman Sachs Global Investment Research estimates AI will expand the total addressable market for automation and enterprise software by roughly 2.5x over the next decade as agentic capabilities extend software's reach into work that was previously human-only.[16] The repricing underway today is happening inside a category that is structurally getting larger, not smaller. The question for incumbents is not survival, but positioning.

AI is now rapidly disrupting software itself, driven by AI-automated code generation and agentic workflows. Unlike the cloud transition, where category leaders provided a clear migration template, the AI transition does not yet offer an established playbook. Few publicly listed software companies have demonstrated a fully AI-native business model at scale. The market does not yet have complete visibility into which incumbents will refound themselves first, and the repricing has been significant.

The iShares Expanded Tech-Software Sector ETF declined approximately 17% in 2026 through June, and about 26% from its October 2025 high, with the sector's 10 largest holdings shedding nearly $800 billion in market capitalization over the same span. Forward P/E multiples compressed from roughly 35x in late 2025 to about 22x, the lowest level since 2014.[17] Regardless of the underlying fundamentals, that uncertainty has driven capital out of enterprise software and into the infrastructure that sits beneath it.

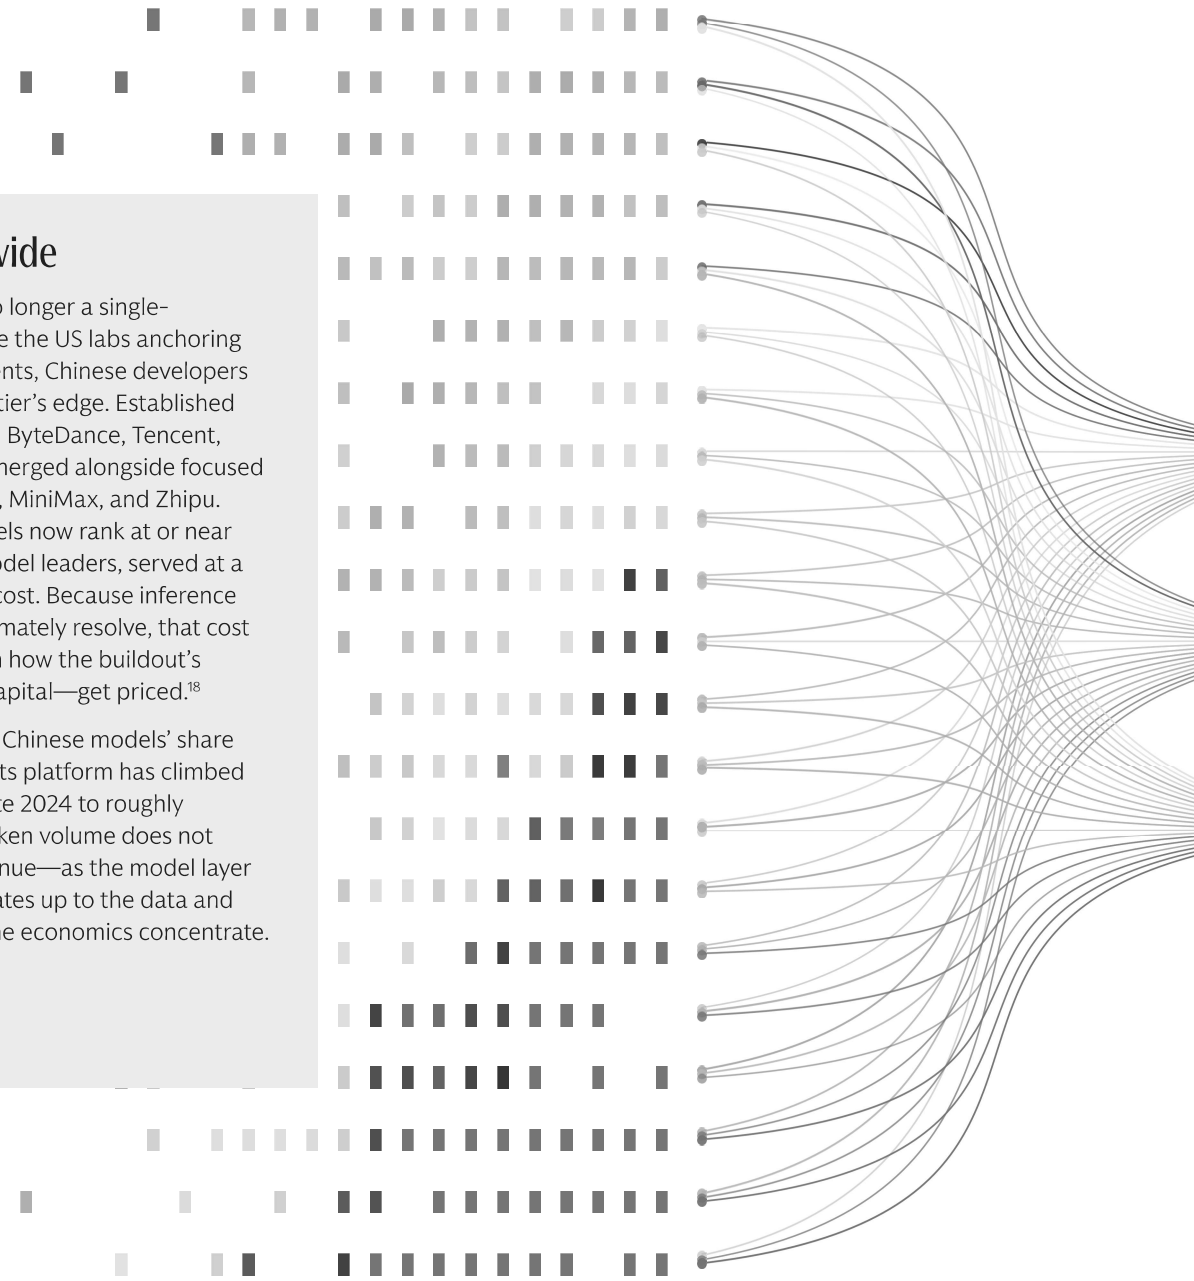
That said, the companies reading this moment as a refounding rather than a threat are pulling ahead. They ship AI capabilities native to their workflows, disrupt their own product economics deliberately, and build toward the emerging value hierarchy rather than defending the existing one.



## Pricing Power Migrates

In the SaaS era, the application layer commanded the premium. Systems organized data, standardized workflows, and monetized access via seats. Pricing power sat with the application layer, which owned the customer relationship. That legacy architecture no longer holds.

Today, premium is migrating toward an outcome-driven architecture, organized around three control points. The outcome layer evaluates software by completed work rather than feature breadth or seat count. The orchestration and agent layer functions as the operating system for enterprise action, governing how work routes and where reliability is maintained. Underpinning both, the data and context layer determines whether AI becomes generic or indispensable: proprietary workflow history and institutional knowledge carry the differentiation as models become commoditized. Training and inference are splitting into different economic categories. Training is the up-front capital event, while inference is the operating cost—running continuously across every user interaction—and where AI economics ultimately resolve. It is also where the US-China cost divide is most visible: Chinese models serve inference at a fraction of US frontier pricing, and the gap has been widening. State-backed providers can price toward cost, where private capital must price to recoup it.



### The China Cost Divide

The model layer itself is no longer a single-geography story. Alongside the US labs anchoring most enterprise deployments, Chinese developers have also pushed the frontier's edge. Established platforms such as Alibaba, ByteDance, Tencent, Baidu, and Xiaomi have emerged alongside focused labs DeepSeek, Moonshot, MiniMax, and Zhipu. Several open-weight models now rank at or near the top of global open-model leaders, served at a fraction of the per-token cost. Because inference is where AI economics ultimately resolve, that cost structure bears directly on how the buildout's compute demand—and capital—get priced.[18]

By OpenRouter's tracking, Chinese models' share of token consumption on its platform has climbed from low single digits in late 2024 to roughly half by early 2026.[19] But token volume does not necessarily equate to revenue—as the model layer commoditizes, value migrates up to the data and outcome layers—where the economics concentrate.

## Cybersecurity as Connective Tissue

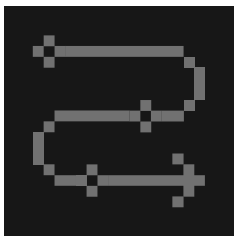
Security has expanded with every computing infrastructure shift—from disconnected systems to the internet, from on-premise to SaaS, and now to AI. Each transition made the "attack surface" broader and the consequences of failure larger, and produced security leaders built specifically for the new architecture. The AI shift is the most expansionary yet. Enterprises must move proprietary data and operational context out of their own environments and into external models, reversing decades of effort to keep that data inside the company's own systems. This shift makes security a precondition for AI deployment rather than a cost line within it, and the market is pricing it accordingly.

The pattern that defined cloud security a decade ago is now visible in AI security at compressed speed. The companies that will dominate AI security in the 2030s are being founded now. For institutions positioning around the AI economy, cybersecurity is not just a vertical; it has become the connective tissue for deployment.

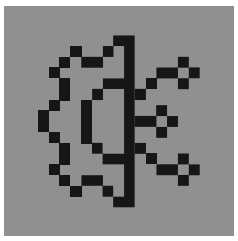
# II. The Industrial AI Software Layer

A parallel restructuring is underway in industrial AI software, the operational layer where incumbents are embedding AI into the tools that design, simulate, and run physical operations. It is less visible than the software sell-off, but the scale is greater, as the overlay is not sector-specific.

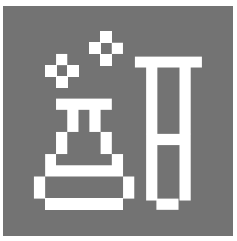
Crucially, this layer is enabling a constant feedback loop where the work product itself (e.g. a car in production) communicates real-time material properties back to the machinery. This enables quality control, design, and process adjustments simultaneously to reduce cost and optimize output.



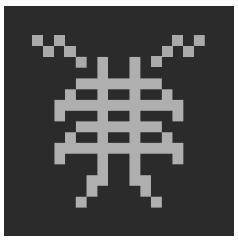
**Manufacturing operators are moving from scheduled maintenance to predictive.**



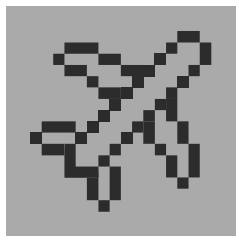
**Automotive engineers train autonomous vehicle models on synthetic data from digital twins.**



**Life sciences are compressing drug discovery and virtualizing trial design.**



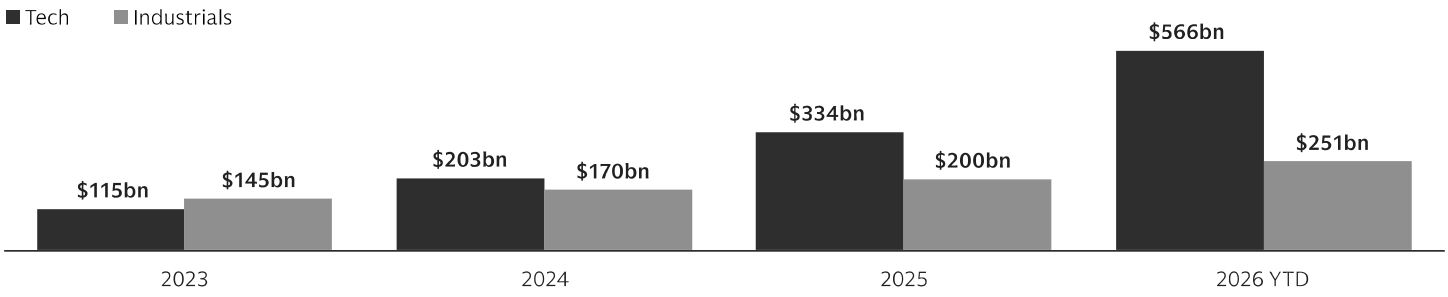
**Energy utilities forecast load and optimize grid dispatch.**



**Aerospace integrates design, testing, and manufacturing into continuous platforms.**

Some of the most consequential platforms are AI-native rather than acquired. Palantir's Foundry has deployed across manufacturing at Airbus, energy at BP, and defense for the US Army's Project TITAN. The category is now standalone, not just an incumbent acquisition target. Major industrial players have spent more than $110 billion acquiring software companies in this space since 2020, and the pace is accelerating.[20]

**Tech and Industrials M&A ($500mm+) continue to see strong YoY momentum**




Source: Dealogic as of June 3, 2026.

## Digital Twins as the Foundation

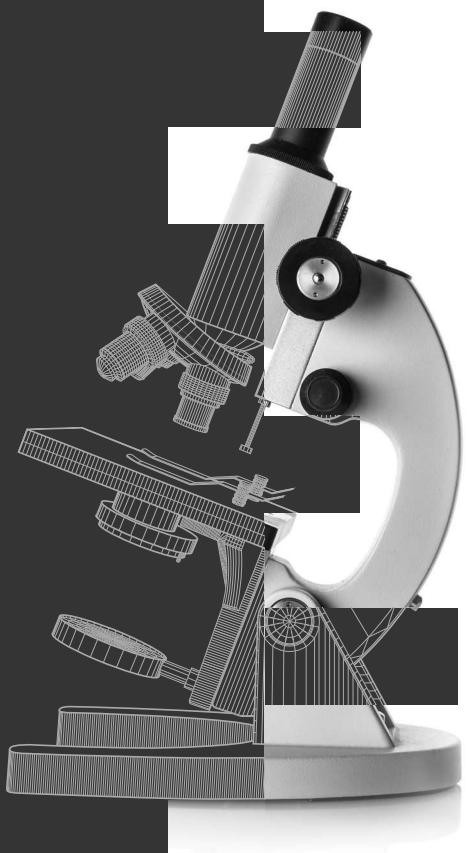
Underpinning this overlay is the digital twin, a virtual rendering of a physical asset, facility, or process built to high fidelity against its real-world counterpart. In 1970, NASA engineers used a ground-based simulator of Apollo 13 to isolate the electrical fault that had ruptured an oxygen tank. Today, industrial operators run persistent digital twins of entire factories, modeling the cost, throughput, and ROI implications of any change before it's made.

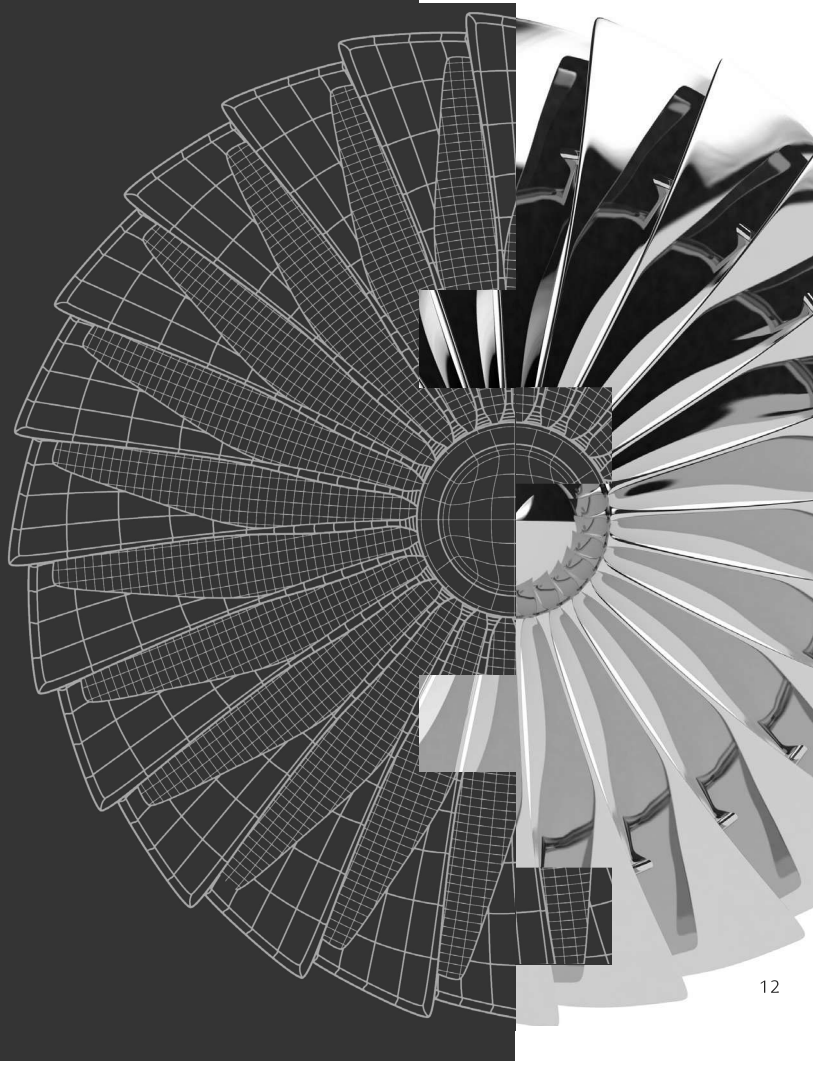
Formula One offers a clear proof of concept. Racing teams run digital twins of their cars, simulating thousands of laps before race day in a risk-free environment, aiming to close the "correlation gap" between simulation and reality.

Siemens' $5.1 billion acquisition of Dotmatics in 2025 sharpens the industrial logic further.[21] Siemens already owned much of the equipment that manufactures pharmaceutical products; with Dotmatics, it moved upstream into designing them, using AI to simulate drug chemistry, compress discovery timelines, and virtualize the trial process itself. The acquisition extends digital twin architecture from factory floor to laboratory bench, and from a single industry into the first end-to-end platform spanning research through manufacturing for life sciences.

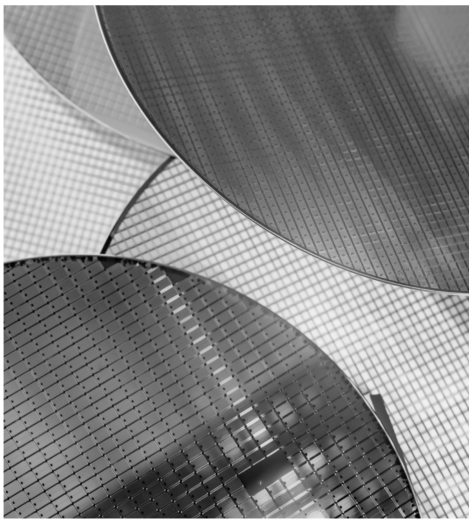
Recent strategic activity illustrates this broader pattern. Synopsys, Inc.'s $35 billion acquisition of Ansys (completed July 2025), the largest pure-software industrial engineering deal of the last decade, consolidated simulation, chip design, and physics-based modeling into a single platform serving semiconductors, aerospace, and automotive.[22] Emerson's roughly $17 billion acquisition of AspenTech, completed in March 2025 with the purchase of the remaining 43% stake, brought process simulation, asset performance management, and industrial AI software into Emerson's automation portfolio.[23]

These platforms are becoming the control plane, and physical operations are reorganizing around them. The M&A consequence: Industrial AI software is now an acquisition target category on par with enterprise data infrastructure, and sector valuation multiples are pricing accordingly. Vertical integration is the unifying logic—each acquirer moving upstream or into an adjacent layer to own the full toolchain within its domain rather than buying scale across it—and it is emerging as an M&A category in its own right. The financing instruments that work for digital infrastructure extend to industrial AI software. The next layer of deployment (humanoids, industrial robots, autonomous fleets) is hardware. The financing architecture for that layer does not yet exist, but it will likely come next.





## The Physical AI Stack—the Brain Layer



Unlike generative models trained on text, physical AI requires multimodal training data (video, force feedback, proprioception), historically scarce and expensive to collect. That scarcity is breaking. The physical AI stack is layered like the software stack, but built from different components: AI brains on Vision-Language-Model and Vision-Language-Action foundations, simulation and digital twin environments generating synthetic training data that real-world testing cannot match at scale, and purpose-built hardware (lightweight actuators, adaptive grippers, high-density batteries, specialized edge compute) that provides machines the physical capability to work alongside humans. Nvidia's open-source Isaac GROOT N1, launched at the 2025 GPU Technology Conference, exemplifies the brain layer.

Around these layers an ecosystem is assembling: chipmakers, cloud providers, software firms, sensor and components, in addition to specialized consultants and integrators. Companies orchestrating multiple layers are positioning to set standards and capture outsized value. Mobileye's early-2026 $900 million acquisition of Mentee Robotics is an example of the cross-layer consolidation the full-stack thesis predicts.[24]

## World Models as the Emerging Foundation

**The three-layer stack now has a name in frontier research: world models.** World models represent cause and consequence in physical and social systems, a capability required for high-value decisions like controlling a robot, managing a supply chain, or coordinating enterprise workflow. Physical world models simulate gravity, friction, thermodynamics, and material behavior, unlocking robotics, autonomous-system, and logistics applications. They are being commercialized at scale today, and require purpose-built simulation environments, synthetic data pipelines, and physics engines.

Social or virtual world models simulate institutional and human behavior, letting firms rehearse competitive moves, stress-test governance structures, and model policy shocks before committing to action. The social variant sits closer to the scenario planning work institutional allocators already do, but its infrastructure is earlier and less mature. Both forms pull on the same compute base.

Consensus forecasts for AI infrastructure (compute, energy, and chips) track transformer-based language models. World models complement language models rather than replacing them. If their trajectory holds, aggregate compute requirements may exceed what current projections anticipate, and the buildout thesis gains a second engine, alongside language models, that the sizing work has not yet priced in.

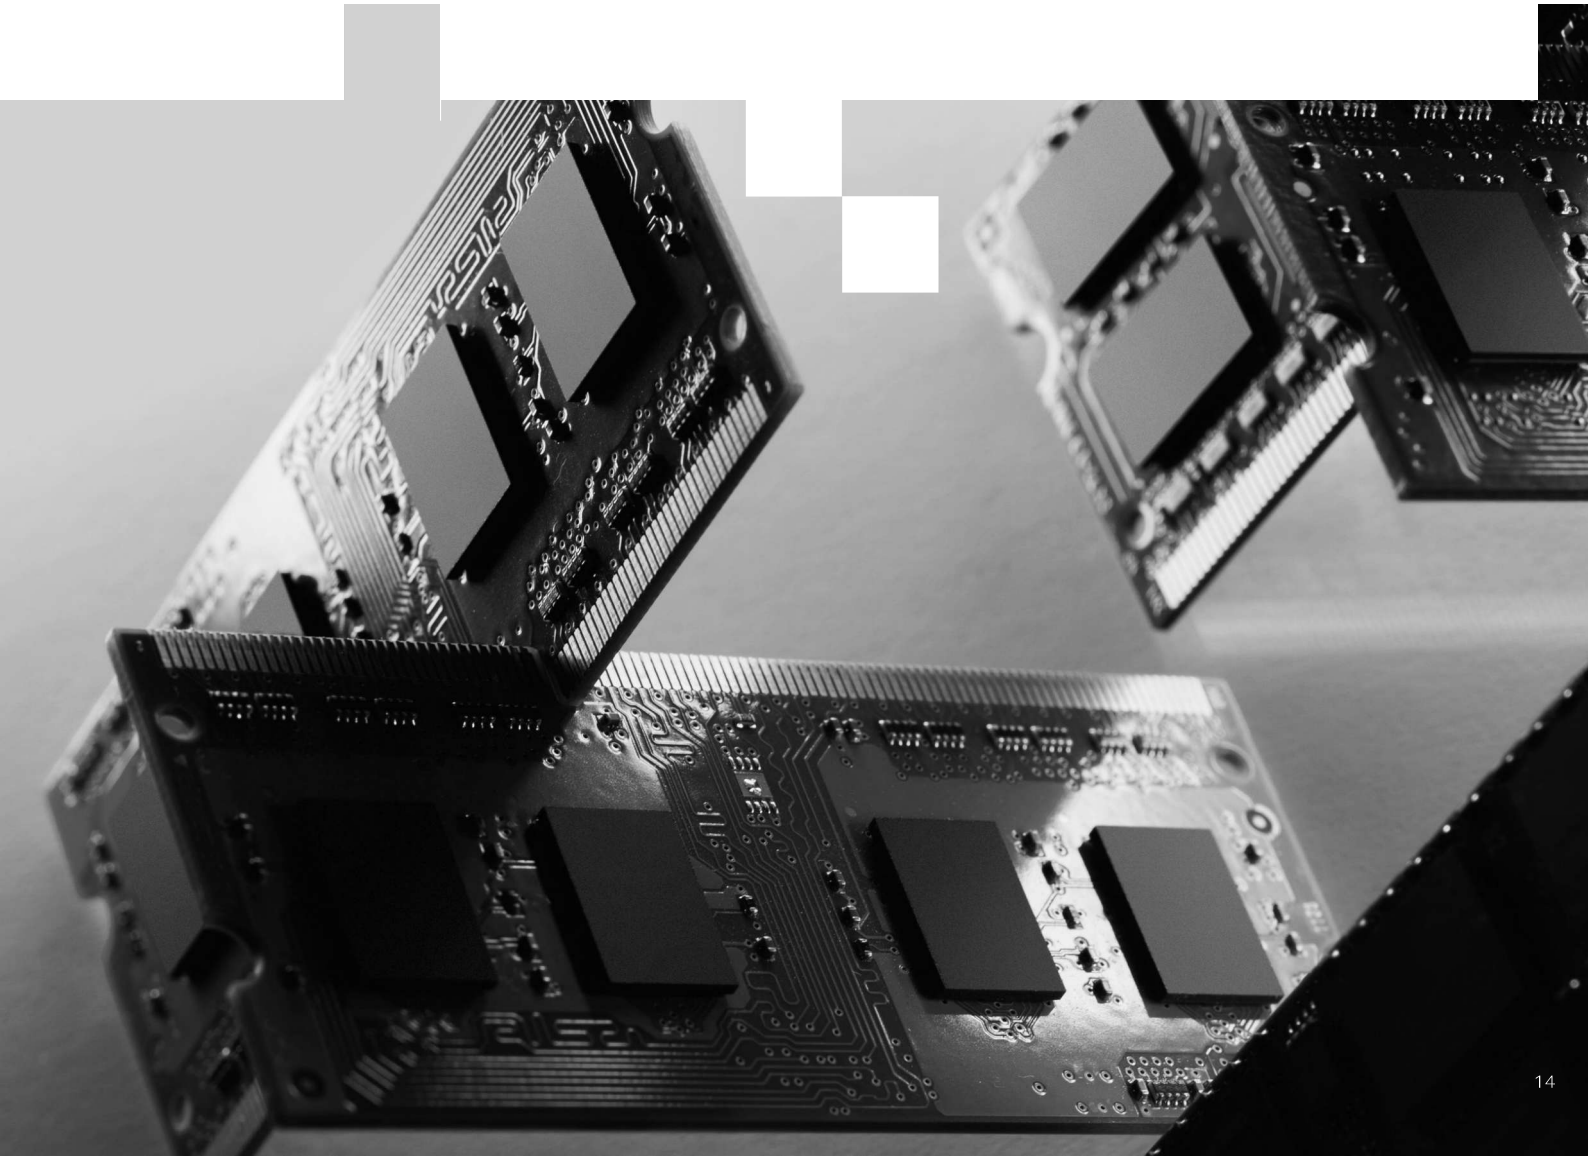


> "After a decade defined by systems that recognize patterns and predict text, the frontier of AI is shifting toward models that understand how the world works—a quiet but decisive change in how machines become intelligent."

**George Lee** | Co-Head of the Goldman Sachs Global Institute

"We are in the middle of an industrial revolution—fueled by the excitement of AI productivity gains and a significant need for investment and CapEx—which presents tremendous opportunities."

**Anthony Gutman** | Co-Chief Executive Officer of Goldman Sachs International and Global Co-Head of Investment Banking



## Five Factors That Will Determine Industrial AI Winners

Industrial AI does not commoditize the way language AI does. Five factors will determine which companies build advantage in the category. The companies that consolidate all five build a durable, competitive position. The companies that build only one or two end up commoditized.

### Physics-based architecture

Unlike language-based AI, Industrial AI cannot scale past the laws of physics. Industrial AI systems must be designed against material properties, thermodynamics, and mechanical tolerances from the start.

### Proprietary data

The best training data for physical systems lives inside manufacturers, utilities, and simulation libraries, not on the public internet.

### Edge deployment capability

Most industrial AI must run inside the machine on the grid node, and on the embedded controller—not just in a cloud data center.

### Certifiability

In aviation, energy infrastructure, and automotive, a single AI hallucination can have catastrophic consequences. Certification is the gating threshold for deployment.

### Workflow integration

The winner embeds AI as a capability layer into existing systems rather than imposing rip-and-replace on decades of operational practice.

# III. Robotics and Physical Deployment

Physical AI refers to intelligent systems that operate in the physical world, including autonomous vehicles and industrial equipment, industrial robots, drones, and humanoid robots. Humanoid robotics dominate the conversation, but non-humanoid systems already operate at far greater scale—and some of the largest bets in the category are not robots at all. Prometheus, cofounded by Jeff Bezos and Vik Bajaj, is building an "artificial general engineer": software meant to design and engineer complex physical systems, from jet engines to drug compounds. It raised $12 billion in June 2026 at a roughly $41 billion valuation, among the largest early-stage investments in the category to date.[25] In addition to venture capital, institutional investors played a significant role, signaling that institutional capital is moving in ahead of commercial deployment economics. AI compresses the invention and engineering cycle itself, scaling industrial transformation and productivity well beyond current AI uses—reshaping processes, industrial design, and automation.

## Humanoids: The Contested Category

The case for humanoids is centered on the need for labor. The United States has 13 million manufacturing workers, more job openings than available workers, and more than one million unfilled materials-handling roles. Goldman Sachs Global Investment Research projects the humanoid market growing from 20,000 units in 2025 to 1.4 million units by 2035.[26] Autonomous industrial equipment such as Deere's field systems and Caterpillar's mining platforms extends the same labor-substitution dynamic into agriculture and heavy industry, where the unit economics are more settled. Deere puts autonomy at the center of its lineup, with machines like the X9 combine as an answer to a worsening labor shortfall, while Caterpillar uses the same approach with its autonomous mining fleet safely hauling over 11 billion tons of material.[27]




"AI optimizes brilliantly within its simulation assumptions, but real materials, tolerances, and the chaos of a live operating environment are never exactly as modeled. A theoretically perfect design can fail because of it."

**Jack Anstey** | Managing Director in Technology, Media, and Telecom Investment Banking

**The commercial model for humanoid robotics is progressing unevenly across three variables that determine viability.**

Capital is the leading indicator—venture funding for companies like Figure, Apptronik, 1X, and Agility has flowed into the space well ahead of revenue, appropriate to the category's stage but underwriting a longer horizon than traditional growth equity models. Hardware is advancing quickly, with multiple platforms demonstrating credible task performance in controlled commercial environments. Workflow integration lags both: embedding humanoids into existing operations, training the workforce, and reaching the unit economics that institutional capital requires all present real challenges. Humanoids need patient, capital-heavy funding suited to long build-and-deploy horizons—closer to deep-tech infrastructure than the faster, capital-light model of venture growth.

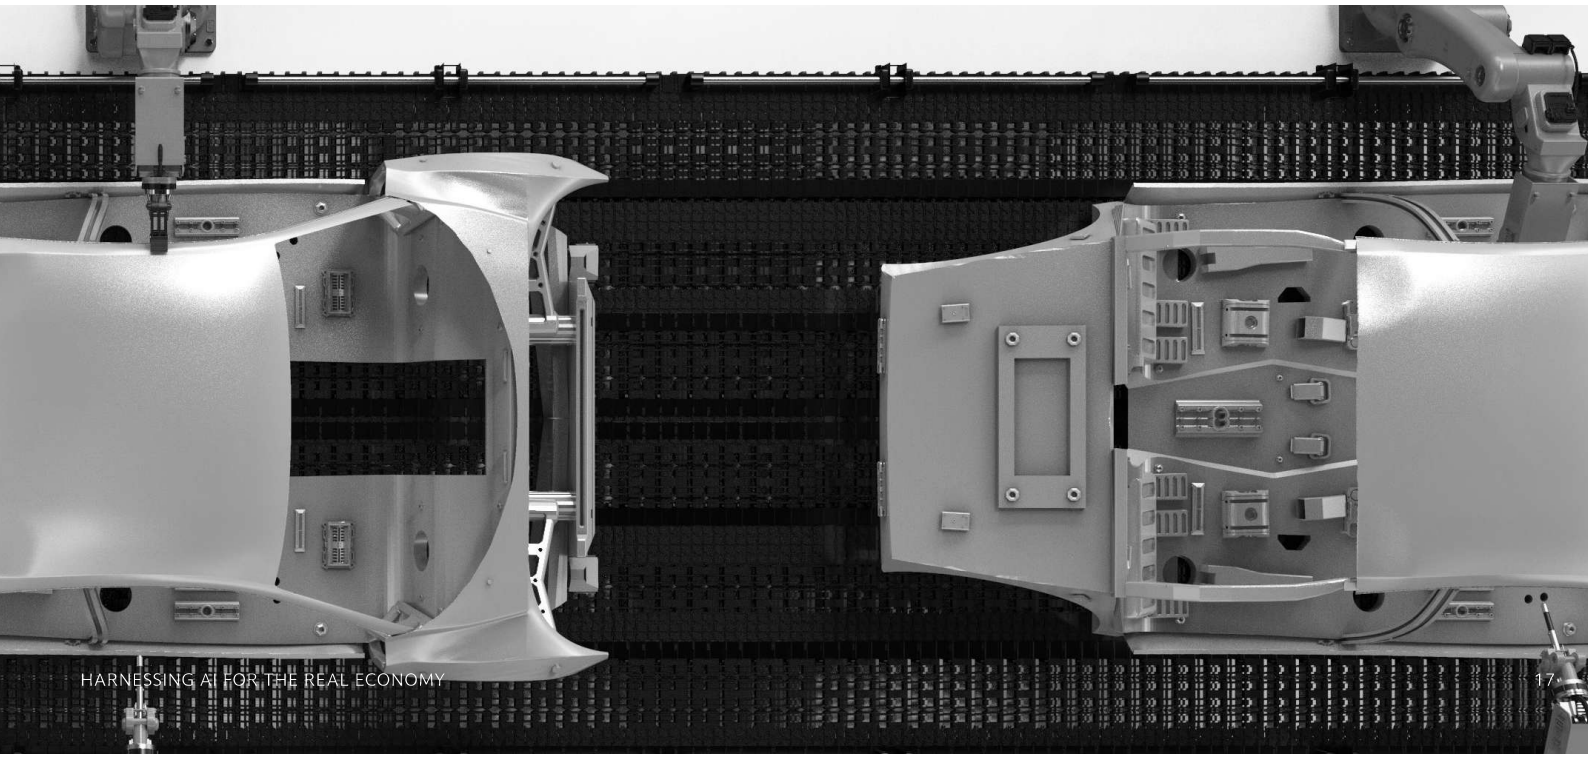
Institutional investors face a significant gap between what these companies aim to build and what can actually be underwritten. Raising the capital that AI-native companies need for compute, hardware, and manufacturing is difficult with few customers and no proven financials. Late-stage private equity is absorbing venture-stage technology risk, with no scaled equivalent of credit wrapping for hardware-intensive deployment and no structure to make digital infrastructure financeable at scale for the physical category.

## Non-Humanoid Robotics Working Today

Today, humanoids are not where industrial robotics generates revenue. Millions of industrial robots, warehouse fleets, and autonomous vehicles are in active use. Next-generation fulfillment centers are now designed around 10-times-higher robotics density from the outset.[28] Symbotic operates dark warehouses for Walmart, as a leading public-market beneficiary of warehouse automation capital flows. In June 2026, Amazon unveiled its latest Proteus autonomous mobile warehouse robot, which responds to natural language. Purpose-built hardware, AI-driven logistics orchestration, and integration into existing retail operations have all attracted investors who understand the unit economics and are putting capital to work.

A cohort of China's industrial AI developers (Unitree, UBTECH, Galbot, AgiBot, Spirit AI, XPENG, LimX) is advancing humanoid/embodied-AI platforms at lower cost. A Goldman Sachs research report on 14 Chinese robotics companies saw single-purpose systems moving toward integrated AI stacks pairing vision-language-action models with world models.[29] The capital behind "data factories" addressing the training-data bottleneck is also scaling: Galbot raised 2.5 billion yuan ($350 million) in early 2026,[30] and Unitree cleared a STAR Market IPO review in June 2026. While cost falls on scale and full-stack control, broad commercial deployment isn't expected until 2027–2029.

Funding is coalescing first around applications whose unit economics are closest to existing underwriting frameworks: logistics automation, predictive maintenance, and warehouse robotics. They are forming unevenly or not at all around applications where venture-stage risk persists. The gap is material, and bridging it has become one of the largest financing questions of the decade.



# IV. Defense as a Strategic Extension

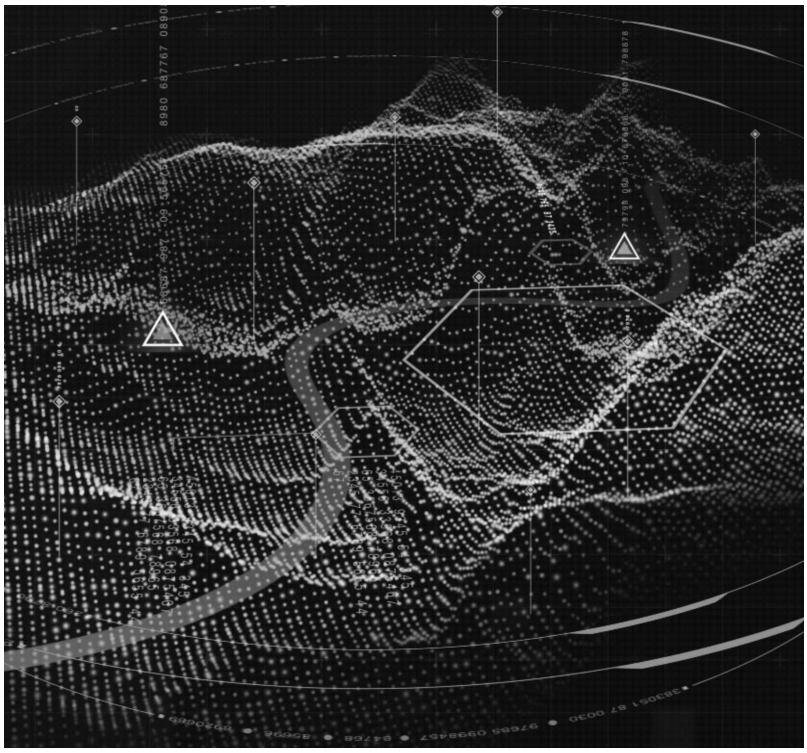
Physical AI deployment, sovereign-government capital partnership, and supply chains repositioned around strategic competition all concentrate in defense. The sector is the most acute test case for the AI economy's industrial logic. Procurement cycles can run a decade or more, capital formation is dominated by a small number of large "primes," and technology adoption moves relatively slowly. AI is changing all three at once, and the shift is visible in companies now challenging the incumbents.

AI in defense is not a replacement for engagement. It is a force multiplier across every stage before the trigger decision, and across the surveillance, communications, and logistics systems that sit around it. Autonomous flight and unmanned aircraft systems are early-mover categories because the human-in-the-loop question is manageable. The operator controls the mission; AI handles flight, navigation, and target identification. The global military drone market reached approximately $20 billion in 2026 and is projected to roughly double over the next decade.[31]

## To Infinity . . .

The operating logic now reshaping defense extends into orbit. Both domains run on the same requirement: AI is not an enhancement but a precondition for operating in today's environment. Both categories face the same constraints—contested signal environments, denied or degraded GPS, latency-sensitive autonomy, and compute that has to function at the edge. The architecture that makes a drone swarm coherent is the architecture that makes satellites operationally useful, and the capital flowing into both has begun to converge.

SpaceX, whose June public market debut saw a record-setting valuation, is a clear example. Following the February 2026 xAI acquisition, it operates as a defense, connectivity, and AI compute business within a single structure. The company is now building an orbital compute layer on top of that network—sun-synchronous-orbit satellites running AI inference at scale, with deployment targeted from 2028. A capital-markets pipeline is forming around space-tech operators whose end customers are defense agencies, and the distinction between a defense listing and a space listing is becoming increasingly narrow.



"As uncomfortable as the idea may be, fully autonomous, AI-executed warfare is an inevitability. It's just a matter of who leads the way."

**Erik Sparks** | Head of Global Resilience Investment Banking

The more consequential change lies in who finances defense technology. Five years ago, defense was a sector many venture capital firms actively avoided. Working with government was considered an exit barrier, a reputational cost, or both. That position has begun to materially shift.

Legacy technology players are now willing to partner on defense. In July 2025, the Department of Defense awarded Google Public Sector and OpenAI contracts valued up to $200 million for military AI, with parallel awards to Anthropic and xAI.[32] The capital markets pipeline is also growing, with defense-focused IPOs in 2025-2026 for Firefly, AEVEX, Hawkeye, York Space Systems, and SpaceX.

M&A in defense, by contrast, remains structurally constrained by a valuation mismatch across the industry and a reticence by challengers to be absorbed into the procurement cultures they were built to disrupt. For now, capital is forming through direct investment and public-market issuance rather than strategic consolidation. A new financing pattern is also emerging—cross-border investment backed by US-government structural support.

## The Challenger Generation



### AMERICAS

Anduril, founded in 2017 with seed capital from Founders Fund, progressed through six venture rounds in eight years, closing a $5 billion Series H in May 2026 at a $61 billion valuation, with a product line spanning air, ground, maritime, and counter-UAS systems unified through the company's software backbone.[33] Its partnership with OpenAI brings frontier AI into national security systems. For institutional investors, the progression signals that defense tech terminal values have moved from speculative to underwritable in under a decade.

### EMEA

A European analog, Helsing, was founded in 2021 and reached a €12 billion valuation in a €600 million Series D in June 2025, deploying AI-powered software and an AI-native strike drone across the German and UK militaries.[34] Much of Europe's strength sits in dual-use engineering rather than pure-play defense. Industrial AI firm PhysicsX raised a $300 million Series C in June 2026 at a roughly $2.4 billion valuation.[35] Its AI-native physics simulation compresses aerospace/defense, semiconductor, automotive, and energy engineering from days to seconds—supplying the defense sector without being defined by it.

### APAC

China presents a structurally different third frame. Domestic defense tech operates inside a military-civil fusion framework that blends state R&D, sovereign capital, and industrial manufacturing in ways US and European systems do not. The procurement pipeline does not face the VC-vs.-government cultural friction that Anduril and Helsing initially navigated, which could create a significant competitive advantage.

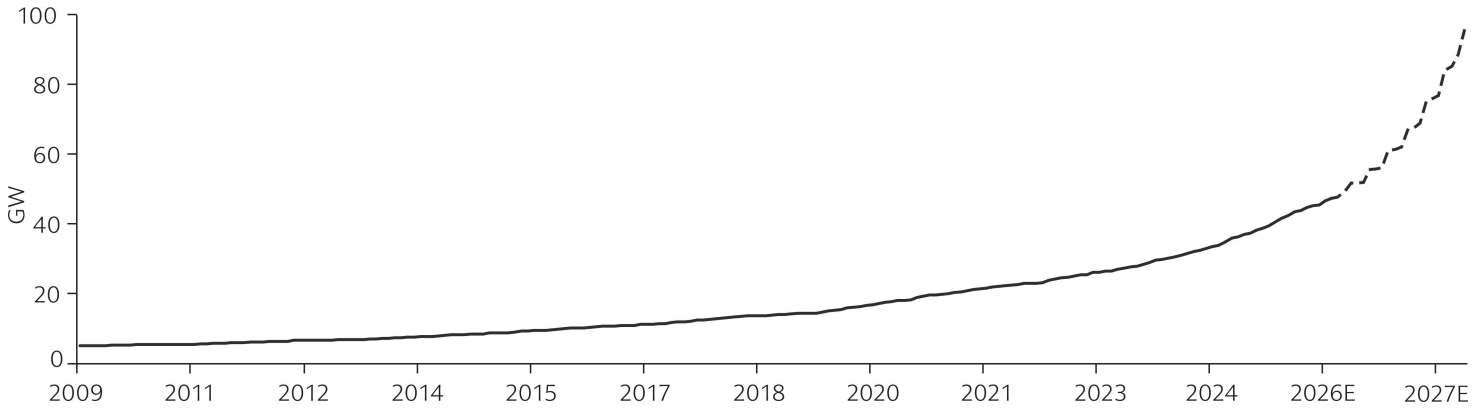
# V. Energy Remains the Binding Constraint

**The capital structures enabling defense, industrial AI, and data center buildout all eventually hit the same gating constraint: power.**

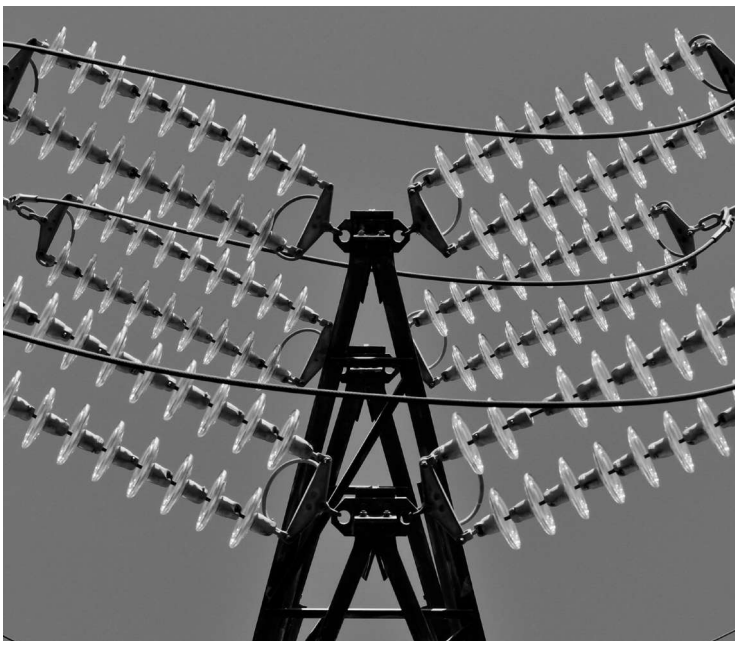
AI training and inference are not the only sources of new load. As autonomous systems proliferate—electric and autonomous vehicles, industrial robots, and sensor-instrumented dark plants—each becomes a continuous generator of data that must be ingested, modeled, and acted on. That processing demand compounds the training and inference load already driving the buildout, and it sits largely outside projections calibrated to today's language-model usage. Physical deployment is not only a consumer of AI; it is a fast-growing source of the compute—and the power—the next phase will require.

Today's grid was built for a different demand profile. Utilities sized capacity around predictable residential and commercial load at roughly 60% average utilization; AI training is unpredictable and clustered in remote markets. A mature AI buildout pushes utilization toward 90%, where most utility operating assumptions cease to hold. The geographic mismatch compounds the problem—Abilene, Texas is now a market whose power demand equals that of San Francisco. Grid interconnection queues in the most important data center markets stretch eight to 12 years, two to three generations of GPU hardware.[36]

**US data center power demand capacity**




Source: Alterio, Goldman Sachs Global Investment Research.



Behind-the-meter power has come into focus as one of a handful of emergent responses. A year ago, fully islanded data centers were rare and not considered a mainstream solution for hyperscale AI infrastructure. Over the last 12 to 18 months, they have moved from an edge case to a credible development pathway driven by grid interconnection delays and the economics of AI compute. Estimates now suggest one-third of future capacity could be islanded, either permanently or as a bridge to eventual grid interconnection. Microsoft, Meta, and Oracle are deploying fully islanded gas generation.

The public markets are now pricing the equipment layer behind this shift: in June 2026, Goldman Sachs acted as joint lead bookrunner on the NASDAQ IPO for INNIO—the gas-engine maker whose Jenbacher systems power islanded data centers—an upsized offering that raised roughly $2.4 billion.[37] Project finance structures originally built for grid-connected power generation are being adapted at scale for behind-the-meter generation dedicated to data center load—a natural evolution in how capital underwrites power infrastructure.

**Nuclear is also back on the table:** Microsoft's Three Mile Island restart agreement (September 2024) and Amazon's small modular reactor investments (October 2024) reflect hyperscalers' evolution from power customers to active participants in power development, and their willingness to bring capital and demand to address power constraints. The uranium conversion bottleneck is a separate concern, with only three Western conversion facilities operating against the rest of global capacity concentrated in Russia and China. Generation is only part of the build. Enabling AI at this scale has also drawn significant capital into the systems around it—electric utility modernization to carry and balance the new load, and water infrastructure to cool it—each a capital-intensive build in its own right. Water is the less-visible side of that build: roughly two-thirds of US data centers built or planned since 2022 sit in water-stressed regions, and data-center water demand is on track to roughly double by 2030—making cooling supply a siting and financing variable, not just an operating cost.[38] These constraints determine whether the AI economy's power layer gets built at the pace the compute layer demands—and whether the new structures bridging utility, hyperscaler, and private capital become the template for capital formation across other physical AI domains.

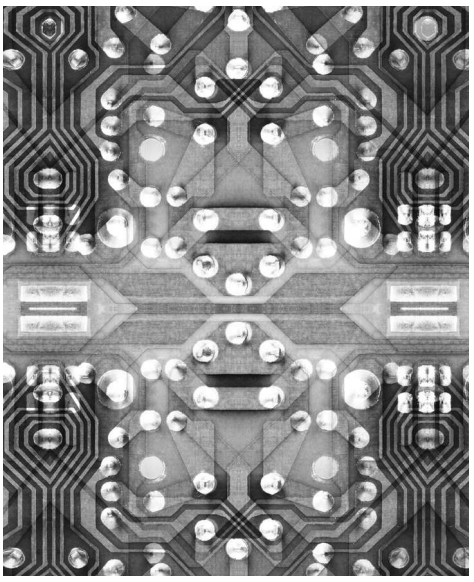
Power is a critical bottleneck—but increasingly, the requisite labor presents a structural constraint of its own. The technical workforce that constructs, wires, cools, and secures this infrastructure is in acute demand, and training cannot happen at the pace capital is being committed. Goldman Sachs Global Investment Research estimates the US power and grid value chain will need more than 500,000 additional workers by 2030, roles that will take three to four years of training to fill. The shortfall is best read as a rate problem: the energy-apprenticeship pipeline ran at 45,000 entrants in 2024 and would need to sustain 65,000 per year now to close it,[39] even before including retirements that outpace new entrants. The implication for capital is direct—financing structures that treat skilled labor as available rather than as a multiyear constraint will encounter delays that capital alone cannot resolve.

These constraints also point beyond the grid. As compute and data-processing demand compounds, the requirement is not only for power but also for moving data at scale with minimal latency—pressure that is already accelerating change within the AI ecosystem, from the physical transmission layer to where computation happens. The same logic that put AI inference into orbit for defense and connectivity extends to energy: processing migrating closer to sustainable, abundant power, including in space. What reads today as a defense and space story is, underneath, the same compute-and-energy constraint expressed in a different domain—and a sign of how far the buildout's architecture may travel from today's earthbound data centers.

**The power and energy sector has a long history of responding to strong economic signals with innovation, investment, and new business models.**

Today's power constraints are real, but so is the industry's ability to adapt. We are seeing innovation emerge across generation technologies, energy storage, financing structures, regulatory frameworks, power market design, and customer contracting models. Just as the shale revolution unlocked vast new energy supply through technological and commercial innovation, the growing demand for power from AI, data centers, and electrification is catalyzing a new wave of solutions. The bottlenecks are significant, but history suggests they are more likely to drive innovation than constrain growth indefinitely.

# VI. The Real Economy's Financing Divide



Across these various sectors, one divide recurs: the financing architecture is mature for some buildout layers and barely formed for others. Physical AI deployment is a visible issue. Humanoids, industrial robotics, and autonomous industrial systems carry venture-stage adoption risk in volumes only long-duration capital pools can currently fund at scale.

For allocator positioning, the implications remain concrete. Digital infrastructure is investable today through familiar instruments—investment-grade bonds, private credit, infrastructure debt funds, real estate trusts. Industrial AI software is accessible through public exposures and selective M&A. Physical AI relies primarily on venture and growth equity, carrying concentration and mismatch risk that institutions must weigh against the upside of underwriting a structurally important category early.

What sits on the other side of this gap is the real economy itself—the manufacturing, energy, robotics, defense, and logistics that AI will significantly change. Overall, the technology is arriving faster than the capital built to fund it. Which sectors thrive and which stall turns less on the tech frontier than on who closes the financing gap—a challenge of capital architecture as much as engineering.



SECTION 03

# Capital Architecture: Financing the Next Phase of Expansion

Strategic partnerships, sovereign capital, and new structured credit instruments are emerging in parallel—drawing on pools of capital that previous tech transitions didn't need to access. Institutions that can structure and direct that capital will shape the next phase of growth.