



Institutional Research Group



Rudy Torrijos
Director, Industry Research
rudy.torrijos@pitchbook.com



Derek Hernandez
Senior Research Analyst, Enterprise SaaS
and Infrastructure SaaS
derek.hernandez@pitchbook.com



Dimitri Zabelin
Senior Research Analyst,
AI and Cybersecurity
dimitri.zabelin@pitchbook.com

Oscar Allaway

Data Analyst

pbinstitutionalresearch@pitchbook.com

Published on April 24, 2026

Contents

Key takeaways	1
Executive summary	4
The rise of digital labor	7
Human capital accelerated	9
AI technology investment stack	15
Tiered AI ecosystem	26
Investment risks	35
Brave new world	39
References	42

EMERGING TECH RESEARCH

Mapping the AI Super-Cycle

Tracking digital labor as it transforms the global economy

Key takeaways

- **New technology investment stack:** Our new 25-layer technology stack maps the flow of investment in a modern AI-centric economy for investors and corporate strategy officers.
- **Digital labor accelerates human capital:** It shifts the individual role from raw technical execution to the high-leverage orchestration of autonomous agent fleets. This digital labor collapses the marginal cost of execution and eliminates administrative duties so that humans can focus on revenue and income growth.
- **Agentic AI pivots the enterprise:** We are at an inflection point where enterprises are moving from AI “pilot projects” to agentic deployments. This shift requires a total re-architecture of the corporate technology stack, as revenue and cost optimization are handed over to autonomous agents executing multistep business goals.
- **The outcome-based revenue:** Enterprise value is shifting from the volume of software licenses sold to the quality of digital work performed. This transition from seat-based ARR to outcome-based consumption models forces a fundamental re-valuation of the software sector, rewarding incumbents who can successfully productize and guarantee the execution of entire high-value workflows.



- **Edge AI and distributed intelligence:** True industrial-grade physical AI cannot occur until expert-level intelligence is fully distributed across the lower tiers of the stack, enabling ultra-low-latency autonomous machines.
- **The human premium:** As expert-level technical execution becomes instantaneous and commoditized, innate human traits—creativity, empathy, flexibility, and trust building—become the only remaining variables for market differentiation. The ultimate killer app of the AI era is the human capable of building customer trust and identifying novel growth opportunities.
- **The rise of individuals as platforms:** We are entering an era where a single human professional scales their impact exponentially by orchestrating fleets of agents. In this paradigm, market value is determined by an individual's orchestration capacity rather than their personal execution speed. This individual-as-a-platform model is a prerequisite to the WaaS economy.
- **WaaS as the new atomic unit:** By 2030, the fundamental unit of economic production will shift from “billable hours” to guaranteed outcomes secured by SLAs. This transforms industries into proactive, recurring revenue engines, enabling higher-margin business models for professional and trade services and higher wages for AI-accelerated employees.
- **Structural inversion of defensibility:** Durable technology moats are aggressively bifurcating to the extremes of the architecture. True defensibility now resides either in the physical computing tiers (owning gigawatt-scale factories, thermal IP, and leading-edge silicon) or the proprietary intelligence tiers (owning customer context and business logic).



- **Sovereign capital as a durable demand layer:** AI infrastructure is no longer purely market-driven; it is increasingly underwritten by long-duration state investment. Governments are active buyers and financiers, treating the technology stack as a national security priority and a prerequisite for technological self-sufficiency.
- **The geopolitical decoupling of the stack:** The AI ecosystem is splitting into geopolitically aligned stacks, anchored by a dominant US-centric architecture and a rapidly developing Chinese alternative. In this scenario, capital flows are determined by political alignments and incentives rather than by market forces. Likewise, supply chains are expanded based on geographical logistics of two unaligned groups rather than optimal service continuity and redundancy for the entire system.
- **The foundation of the next super-cycle:** Over the next decade, market leadership will belong to companies focused on platforms that enable comprehensive SaS digital labor and empower human orchestration. To capture this growth, investors and strategy officers must actively pursue investments across all levels of the stack that balance human-to-human trust building with guaranteed repeatable outcomes. Companies led by management that do not see a future WaaS economy face certain replacement by those that do.



Executive summary

In this note, we launch a new 25-layer technology stack that maps the flow of investment in a modern AI-centric economy. We see work as a service (WaaS) as the new economic unit of production enabled by this stack. In this new paradigm, work will be based on outcome guarantees rather than hours worked. Below, we define every layer of this new stack and describe how digital labor transitions human capital into empowered individual platforms, ultimately changing the nature of work. We offer this as a way for long-term investors and strategy officers to frame current investment and partnership decisions for a five-to-seven-year time horizon.


AI infrastructure spending plans have accelerated dramatically, with hyperscalers and Meta accounting for nearly \$650 billion in 2026, up from \$123.3 billion in 2021. At the start of 2026, enterprises began a 20-year super-cycle to bring digital labor to every part of their organizations. This follows the 20-year investment cycle defined by the deployment of software as a service (SaaS) and cloud services. We anticipate hyperscaler capital expenditure to reach \$1 trillion in annual investments by 2028.

This unprecedented build-out funds the physical infrastructure required to deploy intelligence at scale globally. Hyperscaler capital expenditure, however, pales in comparison to the total enterprise spending required for industries to enable autonomous digital labor across the economy. Gartner estimates that global enterprise IT spending will surpass \$4.7 trillion by 2026.¹ We believe enterprise IT spending will quickly shift to AI, driving innovation up and down the stack.

We segment the AI economy across edge, application, intelligence, infrastructure, datacenter, semiconductor, physical infrastructure, and services. Within enterprise, the transition from SaaS to service as software (SaS) will collapse the marginal cost of executing digital workflows and accelerate the shift to AI digital workforces. SaS describes an outcome-based software product that automates entire employee workflows instead of specific tasks. This shift forces the restructuring of legacy revenue models from seat-based to outcome pricing.



25-layer technology investment stack



Source: PitchBook



Accelerated investments will continue to flow toward the massive physical hardware, power, and thermal management at the foundation of the stack (layers 4 to 13), as high-fidelity model intelligence continuously produces ever-higher productivity rates at ever lower costs.

For incumbent software leaders, we see an equally powerful opportunity to develop a comprehensive, intelligent digital labor force through advancements in physical-layer investments. We see ownership of customer historical data and context-specific business logic as the defining asset at the top of the stack (layers 20 to 25). Without them, a company cannot build durable, defensible moats.

Deploying these new technologies does carry new risks. Endpoint liability, rapid software cannibalization, undefined regulatory headwinds, nonredundant infrastructure, power permitting constraints, and semiconductor supply chain geopolitical risks can undermine the deployment pace of this broader digital economy and fail to meet market expectations.

We see WaaS as the new unit of economic production as a direct result of this new technology investment stack. Individuals and enterprises will begin operating as empowered platforms as early as 2030, providing guaranteed levels of service within consumer, professional/trade, enterprise, and industrial labor sectors. These human-centric platforms will orchestrate fleets of autonomous agents to deliver defined outcomes. Human capital will no longer be defined by the capacity to execute tasks, but by the capacity to manage digital labor.

Despite current market sentiment, digital labor does not displace human capital. It moves humans up a step in the value chain. The killer app of the AI era is the human being. AI fundamentally transforms global industrial organization by re-engineering companies, industries, and economies at the human level. We provide this technology stack for investors and stakeholders to assess their capital opportunities across the new digital-to-physical economy and identify the drivers of growth for this next expansionary super-cycle.



The rise of digital labor

The global infrastructure market is undergoing the largest capital expenditure super-cycle in its history, driven by advanced semiconductor and AI datacenter investment. In our recent notes [SaaS Is Dead, Long Live SaS](#) and [Through the Looking Glass: The Race to Build Enterprise AI](#), we make the case that investors should focus on the SaS transition to digital labor enabled by today's massive AI infrastructure build-out.

We expect that hyperscaler capital expenditure will reach \$1 trillion in annual investments by 2028. This unprecedented build-out will fund the physical infrastructure required to deploy autonomous digital labor at scale. However, this massive forecast pales in comparison to the total enterprise spending over the same period. Gartner estimates that global enterprise IT spending will be \$4.7 trillion in 2026 with a 2024-2029 CAGR of 8.9%.² We believe that over the next five years, enterprise technology infrastructure spending will shift decisively to AI, targeting the infrastructure, intelligence, and applications layers.

Market sentiment is currently paralyzed by a perceived SaaS-pocalypse and fears of hyperscalers' overinvestment in AI capital expenditures. This fear severely discounts the significant economic benefits that users and enterprises are now experiencing from AI. Capital markets have historically rewarded rapidly growing application-layer and asset-light software businesses with high valuation multiples. Long-term customer annual recurring revenue (ARR) has proven to be sticky and reliable across the entire software sector when sold as a service. The introduction of generative AI models and agentic AI applications, however, calls into question the entire value proposition of software, and the survivability of incumbent revenue streams.

We disagree with the current market narrative. The deployment of digital labor through enterprise AI will yield significantly greater revenue opportunities than any other enterprise technology shift. In the intermediate term, reliance on human capital increases rather than decreases as enterprises require humans to supervise and eventually lead teams of digital labor. Employees at all levels will be able to optimize revenue and costs in ways not previously possible. As we saw throughout 2025 in hyperscaler and software vendor earnings reports, end-user employee productivity accelerated with generative AI intelligence gains.



Additionally, the cost-per-watt to supply this higher intelligence will decline significantly through 2028. NVIDIA's next-generation Vera Rubin chip platform is expected to deliver up to 10x more performance per watt for AI inference than the current Blackwell architecture. This will deliver substantially higher productivity for end customers and generate significantly higher-margin potential across the technology investment stack where intelligence is integrated. Companies will likely use margin upside to further invest in digital labor capacity and realignment of business processes to benefit from it.

It is this investment and process realignment that will not only transform business models but also force re-engineering of entire industry value chains. This will function as a sustained productivity catalyst for higher long-term economic growth and reduce market cyclicalities as build-out occurs. We are now seeing this in semiconductor manufacturing, as cyclicalities have all but vanished from five-year revenue forecasts. Design cycles, power improvements, and product introductions accelerate to enable higher-fidelity expert intelligence from large language models (LLMs).

Hyperscalers have jump-started the current investment flywheel, leveraging cash reserves and cash flows from core monopoly advertising, seller fees, and software licenses. Given the dramatic rise in language model fidelity and expert intelligence yielding significant return on investment (ROI) in pilot projects, enterprises have begun to pivot toward digital labor investment in Q1 2026. This was evident in Q4 earnings calls, as highlighted in the [Q4 2025 Advanced Computing PE and VC Trends report](#).

This note provides a modern AI-centric technology stack that maps the entire digital-to-physical economy across 25 technology investment layers. Capital generated at the top of the stack is driven by consumer preferences (layer 22), industrial logic (layer 21), and enterprise workflows (layer 20). This capital flows down to generate sustained ROI for the massive physical infrastructure spending in energy generation, semiconductor manufacturing, and computing capacity (layers 4 to 13).



Human capital accelerated

At the time of this writing, public and private markets are in turmoil. We believe the SaaS-pocalypse was the defining trigger point to pivot enterprises and vendors to agentic AI adoption. Enterprises, investors, incumbents, and startups are now aggressively pushing toward agent-first architectures to address acute investor and customer concerns. As such, we see the need for a new technology investment stack to track this adoption. This allows us and stakeholders to assess growth, competition, and partnerships as we move to digital labor and the AI economy it forms.

By 2030, WaaS will be the new unit of production for early adopting industries, companies, and individuals. We believe this is the primary investment criterion for investors and strategy officers going forward. Below, we define WaaS and how each participant of the labor market is empowered.

Product-market transformation


In a fully realized AI-centric economy, the fundamental unit of economic production transitions from wage hours to guaranteed subscription services. Human capital will be defined by the capacity to orchestrate digital labor to produce outcomes. Individuals and enterprises will become empowered platforms. By directing virtual managers with fleets of specialized agents, humans are freed from administrative tasks. Digital labor can continuously drive down costs as it consolidates these tasks under its purview. This will allow humans to deploy their unique comparative advantages of creativity, empathy, flexibility and human-to-human trust building toward revenue optimization and growth. To understand this new economic landscape, we see three core investment paradigms—SaS, individuals as platforms (IaPs), and WaaS—developing over the next five years, the first of which we introduced in February with our note [SaaS Is Dead, Long Live SaS](#).

Service as software (2026-2031)

Traditional SaaS vendors sell per-seat licenses to enable each subscriber to execute a work task more efficiently. New SaS vendors will sell the completed workflows instead. SaS supplies foundation models, specialized agent frameworks,



Transitional phases to AI-centric services economy



Source: PitchBook



and secure access to proprietary data. This new service will be evaluated on latency, interoperability, and reliability. Vendor revenue models shift from per-user ARR to consumption-based digital labor services. We see successful software incumbents aggressively rearchitecting applications to agentic AI-native architecture with comprehensive workflow capabilities, using hybrid or outcome-based consumption pricing models.

Individuals as platforms (2028-2030)

As SaaS proliferates, the individual professional becomes the chief executive of their own micro-enterprise, regardless of whether they are a W-2 employee, a freelancer, or a trade professional. This framework scales exponentially: One human strategist orchestrates a team of virtual managers, who in turn command hundreds of execution agents. Employees and entrepreneurs transform into IaPs. A person's empathy, creativity, flexibility, and capacity for building trust in relationships dictate the scale of their work. For the IaP, virtual managers translate human intent into workflows, and agents execute at machine speed. We see successful professionals aggressively retraining to become orchestrators, creating their own AI-enabled enterprises that operate independently or within traditional organizations.

While digital execution scales exponentially, human relationship bandwidth remains a high-value scarcity. The IaP model does not scale the human; it scales the human's intent. By leveraging digital labor to handle execution, the professional becomes the provider of understanding. This is the human-to-human context that enables proper handling of client issues, preventing SaaS from becoming a race-to-the-bottom commodity.

Work as a service (2030+)

As digital labor advances through the application tier, work can be priced as guaranteed performance and continuous availability. We define this as WaaS. Within this new paradigm, a buyer purchases a service level agreement (SLA) from an IaP sole proprietor, professional, or trade professional. Instead of paying a one-time invoice, the customer pays a



subscription fee for 24/7 uptime. Likewise, the enterprise relies on the guaranteed performance of its own employee-agent fleets. The employee is paid to orchestrate agents to guarantee performance against KPIs and benchmarks.

In a fully realized AI-centric economy, humans will be tasked with enhancing product-market fit across every industry. Product-market fit is the essential prerequisite for long-term revenue growth and is hyper customer-centric. Understanding the voice of the customer requires human creativity, empathy, flexibility, and trust building, and can only be facilitated by humans. Over the next five years, we expect human capital to be deployed toward revenue generation and away from cost center functions for this reason. Instead, high-fidelity, expert-intelligent systems will autonomously manage large portions of production, administrative, finance, infrastructure, and product development discovery costs.

Labor market transformation

These paradigm shifts fundamentally re-engineer the three major labor markets of the global economy: professionals & trade, enterprise, and industrial. New investment strategies must be able to evaluate how each specific group realizes digital labor.

Professionals & trade

Today's fragmented supply of subcontractors, specialists, and manual laborers will transition into single-point IaP providers, enabling solo operators to scale like midsize franchises. A knowledge professional, such as a lawyer or accountant, will charge a monthly WaaS fee to act as a fractional service provider for dozens of small businesses. Agents will execute contracts and tax filings, while the human provides bespoke strategic counsel and trust. A firm's value will be based on the proprietary data models it builds to train its virtual managers. Similarly, a plumber or electrician will charge a subscription for home infrastructure uptime. Their agents will monitor residential Internet of Things (IoT) sensors, pre-order replacement parts, and schedule preventative maintenance. Humans will only be deployed for the



final physical fix, shifting assurance and liability to the predictive accuracy of the agent fleet. Trades and professions will transition from reactive, hourly-billed enterprises to proactive, subscription-based, recurring revenue services.

Enterprise

The corporation becomes a matrix of IaPs. Chief people officers will transition from hiring for variable technical skills to hiring specifically for agent orchestration capacity. Cost centers will be continuously optimized by digital labor. Agents will handle compliance, accounting, data entry, and baseline customer support with near-zero marginal cost, aggressively stripping out legacy overhead. Freed from operational drag, employees will focus on top-line growth. They will utilize their established trust and personal brand to close complex enterprise deals, their creativity to imagine new product lines, and their experience to pivot strategies. Revenue per employee scales higher as employees use SaS to upskill into IaPs.

The true value of an employee will be driven by the agent-to-human ratio, resulting in a meaningful escalation of enterprise operating cash flow. In deeply competitive markets, production and operating markets will improve, but differentiation will be based on AI-managed product customization and microsegment marketing. By understanding customer intent, companies with highly commoditized products or in low-barrier-to-entry markets will compete on user preferences, not price.

Industrials

The physical world will be managed by agents, greatly enhancing all aspects of a value chain. Robotics will be enabled by intelligence to create zero-downtime environments. Agents will simulate millions of materials science combinations and chemical compounds to accelerate discovery and development, reducing the time from theoretical formulation to physical prototyping from years to days, limited only by physical manipulation constraints and potential resource scarcity.



Virtual managers will oversee global supply chains, autonomously rerouting shipments around weather events or geopolitical friction while optimizing factory floor robotics for maximum yield. Digital twins of physical facilities will be managed entirely by AI, predicting and averting mechanical failures before they happen. Agents will manage the granular, micro-steps of production. Like in the enterprise, industrial workers will transition to revenue-generating functions, including go-to-market, strategic planning, product development, system design, and customer success.

Investment capital transformation

To build highly valuable digital labor, capital will need to flow to companies that enable IaP and not just point solutions. We believe the enterprise AI investment super-cycle begins first with the SaaS infrastructure build-out. We look for technology vendors providing the underlying agent frameworks and virtual manager architectures. These are the suppliers of the new labor force. Early leaders like Salesforce and ServiceNow are successfully transitioning from SaaS seat expansion metrics to comprehensive workflow execution at scale. High net dollar retention will be driven by workflow consolidation and high-fidelity autonomy, not headcount growth. In [Through the Looking Glass: The Race to Build Enterprise AI](#), we offer investors eight AI frameworks to help identify potential market leaders.

Capital must also flow into WaaS platforms and marketplaces. In the long term, we anticipate investments in clearinghouses and platforms that allow consumers and enterprises to easily discover, subscribe to, and manage IaPs. Startups like Fraction, which offers subscription-based engineering outcomes, and Agent.ai, a network for agent-team discovery, are creating the storefronts for this new labor market. These types of platforms will act as the new digital credentialing agencies. The trust, reputation, and SLA reliability of a given IaP are quantified and secured here. The network effect opportunity will create deep moats for the marketplaces that most effectively provide these services.



AI technology investment stack

Our new technology investment stack for the AI economy is segmented into eight groups: edge, application, intelligence, infrastructure, datacenter, semiconductor, engineering, and services across 25 technology investment layers. This fully integrated stack maps the value chain from constructing raw physical infrastructure to deploying intelligent digital labor and, ultimately, to executing autonomous physical work.

Since 2022, hyperscalers and VCs have aggressively funded the semiconductor, datacenter, and intelligence layers of the stack. This investment has been unprecedented, driven by equally unprecedented improvements in end-customer productivity enabled by ever-higher levels of model intelligence. We anticipate annual hyperscaler

AI technology investment stack

Level	Stack tier	Layer	Top five total addressable market (TAM) segments	Top 10 market leaders
25	Edge	Physical AI	Industrial robotics, warehouse automation, commercial autonomous vehicles, humanoids, autonomous ground vehicles	Boston Dynamics, Tesla (Optimus & Semi), Figure AI, Agility, Symbotic, Universal Robots, Fanuc, ABB, Yaskawa, KUKA
24	Edge	AI client	Consumer PCs, smartphones, tablets, smart home electronics, wearable tech	Apple, Samsung, Lenovo, HP, Dell, ASUS, Sony, LG, Google (Pixel), Motorola
23	Edge	Embedded AI	Automotive telematics, industrial IoT, smart appliances, marine & aviation electronics, medical devices	Tesla, Garmin, Oura, Rivian, Samsara, Honeywell, Johnson Controls, Zebra Tech, NXP (Systems), Geotab
22	Application	Consumer AI	Social media, video streaming, digital music, mobile gaming, e-commerce	YouTube, Shopify, Spotify, TikTok, Netflix, Roblox, Airbnb, Duolingo, Stripe, Amazon
21	Application	Industrial AI	Defense & aerospace tech, heavy manufacturing systems, global supply chain logistics, healthcare IT, core banking & fintech	Veeva Systems, Palantir, Epic Systems, Procore, Guidewire, Tyler Tech, nCino, Amdocs, Bentley Systems, Constellation
20	Application	Enterprise AI	Customer relationship management (CRM) and sales tech, enterprise resource planning (ERP) & supply chain software, HRIS & payroll, financial accounting, analytics, knowledge management, governance risk & compliance	Salesforce, SAP, Workday, Oracle, Microsoft, Adobe, ServiceNow, Intuit, Atlassian

Source: PitchBook



capital expenditure reaching \$1 trillion by 2028. We also expect capital expenditure growth to continue through 2030, albeit at a slower pace than from 2024 to 2026. Please see our hyperscaler TAM forecast in the [Q3 2025 Launch Report: Advanced Computing](#).

We are now at the next inflection point of growth: the rise of the enterprise AI. This was evident in Q1 public company earnings commentaries, as customers are moving from AI pilots to agentic deployments. This pivot requires a new technology investment stack to track the rise of digital labor investment across various industrial sectors. As enterprises rearchitect their technology stacks, additional margins generated by revenue and cost optimization will drive profound changes in human capital productivity and capital expenditure. For layers 20 to 22, the application tier, we expanded the stack to include important subgroups that are direct beneficiaries of agentic digital labor.

AI technology investment stack (continued)

Level	Stack tier	Layer	Top five TAM segments	Top 10 market leaders
19	Intelligence	Agentic AI	Robotic process automation, business process management software, autonomous agents, agent operations (AgentOps), AI trust, risk & security management (AI TRiSM)	Microsoft, OpenAI, Anthropic, DeepMind, Cognition, CrewAI, LangChain, LlamaIndex, Salesforce, AgentOps
18	Intelligence	Generative AI	Foundation models, LLM APIs, multimodal generation, machine learning (MLOps) and LLM operations (LLMOps), code copilots	OpenAI, Anthropic, Google, Meta, Mistral AI, Cohere, xAI, AI21 Labs, Databricks, NVIDIA
17	Intelligence	Data fabric	Relational databases, data warehouses, data lakes, vector databases, data governance, master data management (MDM)	Oracle, Microsoft (SQL Server), Snowflake, Databricks, AWS (Redshift), IBM (Db2), Collibra, MongoDB, Pinecone, Neo4j
16	Infrastructure	Orchestration	Container management, API management, CloudOps and observability, service mesh, DevOps, DevSecOps	Google (GKE), Red Hat, Microsoft (AKS), AWS (EKS), F5, MuleSoft, Kong, HashiCorp, Docker, Apigee
15	Infrastructure	Virtualization	Server virtualization, hypervisors, desktop virtualization, storage virtualization, SDN	VMware, Microsoft (Hyper-V), Red Hat, Canonical, Citrix, Nutanix, Oracle, IBM, SUSE, Amazon (Nitro)
14	Infrastructure	Secure delivery	Content delivery networks, WAF, load balancing, DDoS protection, DNS services	Cisco, Arista, F5, Palo Alto, Fortinet, Cloudflare, Akamai, Juniper (HPE), Fastly, Check Point

Source: PitchBook



A new stack was also needed to define the relationship of edge devices to the full technology stack. This tier includes autonomous machines, AI clients, and embedded AI. True intelligence for machine automation and human assistance cannot be achieved until all levels of the stack are enabled with high-fidelity agentic capabilities. Both PCs and mobile phones achieved maximum penetration when PCs became fast and networked (wifi) and mobile phones became fast and smart (5G, iOS/Android). In the same way, we see ultra-low-latency expert-level intelligence as an essential prerequisite for the rise of edge AI.

We expanded the application and semiconductor tiers to define the complex ecosystem needed to bring digital labor to all verticals of the economy. The top five TAM segments and top 10 market leaders for each technology investment layer are included.

AI technology investment stack (continued)

Level	Stack tier	Layer	Top five TAM segments	Top 10 market leaders
13	Datacenter	Hyperscale cloud	IaaS, PaaS, cloud storage, GPU as a service, cloud compute	AWS, Azure, Google Cloud, Oracle Cloud, Alibaba, IBM Cloud, Tencent, CoreWeave, Lambda, Baidu
12	Datacenter	Colocation & real estate	Retail colocation, wholesale datacenters, datacenter REITs, managed hosting, edge datacenters	Equinix, Digital Realty, CyrusOne, Vantage, QTS, Switch, CoreSite, Iron Mountain, NTT, DataBank
11	Datacenter	IT hardware	Enterprise servers, storage arrays, mainframes, network switches, enterprise routers	Dell, HPE, Supermicro, Lenovo, Cisco, Inspur, Quanta, Foxconn, Wistron, IBM
10	Datacenter	Power & thermal	UPS systems, precision cooling & computer room air conditioners, liquid cooling, power distribution units, backup generators	Vertiv, Schneider Electric, Eaton, Delta, Legrand, Stulz, Rittal, Munters, CoolIT, Nortek
9	Semiconductor	Quantum technologies	Superconducting qubits, trapped ion qubits, quantum software, quantum annealing, photonic QC	IBM, Google, Microsoft, Quantinuum, IonQ, Rigetti, D-Wave, PsiQuantum, Infleqtion, Pasqal
8	Semiconductor	Chip design	Fabless semiconductor design, GPU architecture, CPU architecture, ASICs and FPGAs, EDA software	NVIDIA, AMD, Intel, Broadcom, Qualcomm, ARM, Marvell, MediaTek, Apple, Cerebras
7	Semiconductor	Semiconductor supply chain	Semiconductor foundries, lithography equipment, wafer manufacturing, IC packaging and testing, electronic materials	TSMC, ASML, Applied Materials, Lam Research, KLA, Tokyo Electron, Samsung, GlobalFoundries, Synopsys, Cadence

Source: PitchBook



AI technology investment stack (continued)

Level	Stack tier	Layer	Top five TAM segments	Top 10 market leaders
6	Physical infrastructure	Manufacturing infrastructure	Electronics manufacturing services, original design manufacturers, electronic component distribution, industrial tooling, factory machinery	Jabil, Flex, Sanmina, Fanuc, Keyence, Arrow Electronics, TD Synnex, Omron, Celestica, WPG Holdings
5	Physical infrastructure	Communication infrastructure	Telecom base stations, fiber-optic networks, cell towers, subsea cables, satellite communications	Ericsson, Nokia, Huawei, Cisco, Corning, CommScope, Crown Castle, American Tower, Ciena, ZTE
4	Physical infrastructure	Power infrastructure	Baseload generation (nuclear, coal, and gas), renewable energy, high-voltage transmission, smart grids, energy storage	NextEra Energy, Duke Energy, EDF, State Grid, Constellation, Iberdrola, Southern Co, Enel, Orsted, TVA
3	Services	IT services	Value-added resellers, managed service providers, IT hardware deployment, cloud migration, network integration	CDW, Insight Enterprises, SHI International, WWT, Presidio, Softchoice, ePlus, Computacenter, NTT, Kyndryl
2	Services	Engineering & construction	Heavy civil construction, datacenter construction, industrial plant construction, infrastructure engineering, construction management	Bechtel, Fluor, Jacobs, Kiewit, Turner Construction, Skanska, PCL, Vinci, Bouygues, Hochtief
1	Services	Professional services	Management consulting, systems integration, IT strategy advisory, audit & accounting, business process outsourcing	Accenture, Deloitte, IBM Consulting, TCS, Capgemini, Infosys, Wipro, Cognizant, DXC, NTT DATA

Source: PitchBook



Application investment stack

The application tier (layers 20 to 22) is the intelligence demand engine of the AI economy. It serves as the interface through which human intent is translated into digital labor execution. We segment this tier into consumer, industrial, and enterprise workflows, with further segmentation to provide stakeholders with a more granular TAM market mapping. The top five TAM segments and top 10 market leaders for each technology investment sublayer are included.

Application investment stack

Stack layer	Stack industry number	Economic focus	Targeted end markets	Market categories	Market leaders
Consumer AI	22.01	Consumer economy	Lifestyle & social	Social media, video streaming, gaming, digital music, creator economy	YouTube, Meta, Netflix, Roblox, Epic Games, Spotify, TikTok, Tencent, Twitch, Patreon
Consumer AI	22.02	Consumer economy	Personal services	Edtech (B2C), telehealth, fitness tech, mental wellness, skill building	Duolingo, Teladoc, Hims & Hers, Peloton, Strava, Calm, BetterUp, Coursera, Babbel, MasterClass
Consumer AI	22.03	Consumer economy	B2C commerce & retail	E-commerce, point-of-sale, online travel, food delivery, consumer fintech	Shopify, Amazon, Toast, Block (Square), Airbnb, Expedia, DoorDash, Uber, Stripe, PayPal
Industrial AI	21.01	Knowledge economy	Professional services	Legal tech, AI legal & audit, tax & accounting, SMB ERP, insurance (P&C)	Relativity, Clio, Harvey, Ironclad, Vertex, Avalara, Intuit, Xero, Guidewire, Duck Creek
Industrial AI	21.02	Knowledge economy	Institutional fintech	Wealth & asset management, capital markets, commercial banks, RWA, custody	Addepar, Envestnet, Broadridge, Clearwater Analytics, nCino, Temenos, Securitize, Ondo, Fireblocks, Anchorage



Application investment stack (continued)

Stack layer	Stack industry number	Economic focus	Targeted end markets	Market categories	Market leaders
Industrial AI	21.03	Knowledge economy	Healthcare	Hospitals (EHRs), clinical care, pharmacy tech, health insurtech, oncology	Epic Systems, Oracle Health, Athenahealth, Doximity, GoodRx, DrFirst, Clover Health, Oscar, Flatiron, Tempus, ModMed, Abridge, Ambience, OpenEvidence, Doctronic, Lumeris, Amazon Health
Industrial AI	21.04	Knowledge economy	Life sciences	Pharma research & development, clinical trials, genomics, biomanufacturing, regulatory	Eli Lilly, Novartis, Pfizer, Xaira Therapeutics, Generate:Biomedicines, insitro, Formation Bio, Isomorphic Labs, Caris Life Sciences, Illumina, Schrödinger, Veeva Systems, Medidata, Benchling, Dotmatics, DNAnexus, Apprentice, MasterControl, Certara
Industrial AI	21.05	Asset economy	Real estate & proptech	Commercial real estate, construction tech, property management, facilities management, real estate finance	CoStar, Procore, AppFolio, VTS, Yardi, Bentley Systems, Trimble, Entrata, MRI Software, Autodesk
Industrial AI	21.06	Asset economy	Industrial & heavy manufacturing	Heavy industries, factory auto, PLM, digital twins, process manufacturing	Siemens (digital), Dassault Systèmes, Rockwell, PTC, AVEVA, Hexagon, Cognite, Tulip, Plex, Emerson
Industrial AI	21.07	Asset economy	Supply chain, logistics & transportation	Commercial transportation, fleet management, WMS, freight forwarding, last mile	Manhattan Associates, Blue Yonder, Samsara, Motive, Flexport, Descartes, Project44, Amadeus, Sabre, Locus



Application investment stack (continued)

Stack layer	Stack industry number	Economic focus	Targeted end markets	Market categories	Market leaders
Industrial AI	21.08	Asset economy	Agriculture & natural resources	Agtech, mining tech, basic materials, chemicals, forestry	John Deere, Hexagon (mining), Trimble, AspenTech, Climate Corp, Farmers Edge, Seeq, Kespry, MineSense, Datamine
Industrial AI	21.09	Asset economy	Energy & utilities	Grid management, renewables, oil & gas, water utilities, carbon & ESG	GE Vernova, Siemens (energy), Uplight, Stem, Schlumberger, Enverus, Innovyze, Watershed, Persefoni, Itron
Industrial AI	21.10	Asset economy	Sovereign & public sector	Government admin, civic tech, public safety, transit authorities, tax & treasury	Tyler Tech, Accela, Granicus, OpenGov, Axon, Motorola Solutions, CentralSquare, Mark43, CivicPlus, PayIt
Industrial AI	21.11	Asset economy	Aerospace, defense & space	Defense tech, intelligence & C2, orbital networks, autonomous weapons, aerospace	Palantir, Anduril, SpaceX (Starshield), Shield AI, Epirus, Rebellion Defense, HawkEye 360, Skydio, BAE (digital), Lockheed (cyber)



Application investment stack (continued)

Stack layer	Stack industry number	Economic focus	Targeted end markets	Market categories	Market leaders
Enterprise AI	20.01	Revenue optimization	Growth	CRM, marketing automation, sales enablement, CPQ, partner relationship management	Salesforce, HubSpot, Adobe, Oracle, Microsoft, SAP, Zoho, ActiveCampaign, Braze, Klaviyo
Enterprise AI	20.02	Revenue optimization	Service	CX, IT service management, field service management, contact center as a service, knowledge management	ServiceNow, Zendesk, Freshworks, Salesforce, Atlassian, Intercom, Genesys, Five9, BMC, Ivanti
Enterprise AI	20.03	Revenue optimization	Creative	PLM, CAD, version control, design collaboration, agile project management	Atlassian, Microsoft (GitHub), GitLab, Adobe, Autodesk, Siemens, Dassault Systèmes, Synopsys, JetBrains, Figma
Enterprise AI	20.04	Cost optimization	Operations	ERP, SCM, inventory management, procurement, order management systems, analytics, knowledge management	SAP, Oracle, Microsoft, Infor, Epicor, Blue Yonder, Manhattan Associates, Kinaxis, Odoo, IFS
Enterprise AI	20.05	Cost optimization	People	Human capital management, payroll processing, applicant tracking systems, workforce management, benefits admin	Workday, ADP, UKG, Paycom, Paylocity, SAP, Oracle, BambooHR, Dayforce, Gusto
Enterprise AI	20.06	Cost optimization	Finance	Financial planning & analysis, general ledger, AP/AR, treasury management, tax compliance	Oracle (NetSuite), Intuit, Sage, Xero, Coupa, SAP (Concur), Workday, BlackLine, Expensify, Bill
Enterprise AI	20.07	Cost optimization	Governance	GRC, enterprise risk management, audit management, regulatory change management, IT security & risk	ServiceNow, OneTrust, Archer, AuditBoard, MetricStream, IBM, SAP, Diligent, LogicGate, Navex



Semiconductor investment stack

The semiconductor tier (layers 7 to 9) is the physical foundation of the digital labor economy. We segment this tier across the semiconductor supply chain, chip design, and quantum technologies to map the flow of infrastructure capital. Investors can use this taxonomy to identify highly defensible, capital-intensive value chain vendors that directly accelerate agentic, generative, and embedded intelligence in their higher tiers. Cost/watt limitations gate digital labor adoption. As a result, the semiconductor market is the primary beneficiary of hyperscaler capital expenditure as design and leading-node capacity availability cycles quickly shorten. We include the top five TAM segments and top 10 market leaders for each technology investment sublayer.

Semiconductor investment stack

Stack layer	Stack industry number	Economic focus	Targeted end markets	Market categories	Market leaders
Quantum computing	9.01	Quantum technologies	Quantum hardware	Superconducting qubits, trapped ion, neutral atoms, photonics, quantum annealing	IBM, Google, Quantinuum, IonQ, D-Wave, Rigetti, Pasqal, PsiQuantum, QuEra, Silicon Quantum Computing
Quantum computing	9.02	Quantum technologies	Quantum software & cloud	Quantum as a service, algorithm dev, error correction, quantum ML, optimization	Microsoft, AWS, IBM, SandboxAQ, Classiq, Q-CTRL, Riverlane, Zapata, Strangeworks, QC Ware
Chip design	8.01	Semiconductor chip design	GPU/CPU/ ASSP/ASIC	Datacenter GPUs, enterprise CPUs, custom AI silicon, networking ASICs, edge processors	NVIDIA, AMD, Intel, Broadcom, ARM, Qualcomm, Cerebras, Groq, Tenstorrent, SambaNova
Chip design	8.02	Semiconductor chip design	Memory (DRAM & NAND)	High bandwidth memory, datacenter DDR5, enterprise SSDs, mobile LPDDR, storage class memory	Samsung, SK Hynix, Micron, Western Digital, Kioxia, CXMT, YMTC, Solidigm, Winbond, Nanya
Chip design	8.03	Semiconductor chip design	Analog, mixed signal & MCU	Automotive MCUs, industrial control, RF ICs, data converters (ADC/DAC), amplifiers	Texas Instruments, NXP, STMicroelectronics, Analog Devices, Microchip, Renesas, Onsemi, Wiliot, Ayar Labs, Black Semiconductor
Chip design	8.04	Semiconductor chip design	Power management IC	Datacenter power stages, EV battery management, smart grid control, consumer fast charging, voltage regulators	Infineon, Texas Instruments, STMicroelectronics, Analog Devices, Wolfspeed, NXP, Cambridge GaN Devices, Egghtronic, Silanna Semiconductor, Navitas
Chip design	8.05	Semiconductor chip design	Optoelectronics	Silicon photonics, optical transceivers, image sensors (CMOS), LEDs/lasers, LiDAR components	Sony, Broadcom, Lumentum, Coherent, OmniVision, Ayar Labs, Lightmatter, Celestial AI, Saliency Labs, Xsight Labs



Semiconductor investment stack (continued)

Stack layer	Stack industry number	Economic focus	Targeted end markets	Market categories	Market leaders
Chip design	8.06	Semiconductor chip design	Sensors, discretes & passives	MEMS sensors, power MOSFETs, IGBTs, ceramic capacitors, precision resistors	Bosch, STMicroelectronics, Infineon, TDK, Murata, Yageo, Uhnder, Metawave, Prophesee, Vayyar
Supply chain	7.01	Semiconductor supply chain	Semiconductor materials & wafers	Bare silicon wafers, silicon carbide (SiC), gallium nitride (GaN), photoresists, electronic gases	Shin-Etsu, SUMCO, GlobalWafers, Siltronic, Entegris, JSR, Heraeus, Brewer Science, EMD Electronics, Nanosys
Supply chain	7.02	Semiconductor supply chain	EDA & IP	Logic synthesis, physical verification, core architecture IP (ARM/RISC-V), interface IP, emulation	Synopsys, Cadence, Siemens EDA, ARM, Rambus, SiFive, Codaip, Achronix, Ascenium, Movellus
Supply chain	7.03	Semiconductor supply chain	Foundry	Leading edge (sub-5nm), mature nodes, specialty analog/RF, SiC/GaN foundry, advanced memory foundry	TSMC, Samsung Foundry, Intel Foundry, GlobalFoundries, UMC, SMIC, Rapidus, SkyWater, Nexchip, DB HiTek
Supply chain	7.04	Semiconductor supply chain	Wafer fabrication equipment	EUV/DUV lithography, deposition (CVD/ALD), plasma etch, metrology & inspection, ion implantation	ASML, Applied Materials, Lam Research, Tokyo Electron, KLA, Screen, ASM International, ZEISS SMT, EV Group, Canon



Semiconductor investment stack (continued)

Stack layer	Stack industry number	Economic focus	Targeted end markets	Market categories	Market leaders
Supply chain	7.05	Semiconductor supply chain	Advanced packaging	2.5D/3D IC, CoWoS, chiplets, fan-out wafer level packaging, heterogeneous integration	TSMC, ASE Group, Amkor, Intel, Samsung, JCET, Deca Technologies, Promex Industries, NHK Spring, Shinko
Supply chain	7.06	Semiconductor supply chain	Outsourced assembly & test	Wire bonding, flip chip, wafer testing, final testing, system-in-package	ASE Group, Amkor, JCET, Powertech, Tongfu Microelectronics, Huatian, UTAC, Carsem, Unisem, King Yuan



Tiered AI ecosystem

Agentic AI did not exist two years ago, and it is now actively forcing the transformation of every existing software function and feature. This is causing a significant reassessment of what business defensibility really means. In the physical layers, owning leading-edge designs, manufacturing, and computing services with the fastest time-to-market is creating the highest barriers to entry (AI semiconductors, AI factories, and power & thermal management). In the intelligence and software layers, owning and orchestrating trusted customer data, business logic, and entire deployment workflows will create the strongest moats.


Over time, LLMs will become commoditized as high-fidelity expert intelligence plateaus across all vendors. Digital labor, however, can maintain long-term business viability so long as customer data and business logic remain proprietary. We see LLM companies with comprehensive agentic enterprise solutions and enterprise software companies with demonstrated agentic functionality becoming the major providers of digital labor. Over the next five years, the initial build-out of enterprise AI should enable the long-anticipated industrial adoption of edge AI as high-fidelity intelligence is distributed throughout the stack.

We provide detailed descriptions of our new technology investment stack below. Our intent is to provide investors and strategy officers with a new market segmentation to identify competitive threats and potential partnerships. The following descriptions attempt to define the technical function, structural dependencies, and economic boundaries of each layer. It provides an eight-segment, mutually exclusive framework that separates services, physical infrastructure, semiconductors, datacenters, infrastructure software, intelligence, applications, and edge. This provides a new investment industry segmentation for understanding the modern digital-to-physical AI stack.

This segment represents the physical embodiment of AI architecture. It encompasses endpoint hardware with embedded/bundled software that translates cognitive intelligence into real-world actions. These are devices in which digital labor becomes autonomous physical actions or directly assists human users.



Technology investment stack tier:
Edge AI




Edge AI tier

- **Layer 25—physical AI:** The embodiment of anthropomorphic execution. The robotic endpoint is achieved when autonomous intelligence is deployed in commercial machinery, humanoids, and industrial automation. It functions as the physical hands and legs of the AI architecture, translating digital logic into kinetic action across factory floors, supply chains, retail, and the home.
- **Layer 24—AI client:** The direct human-compute interface. This encompasses consumer and enterprise endpoint devices, including AI PCs, smartphones, tablets, and future interface devices. It serves as the primary gateway to localized processing, allowing humans to physically interact with the digital labor force with reduced cloud dependency and latency.
- **Layer 23—embedded AI:** This consists of small language models running locally on industrial sensors, smart vehicles, and IoT devices. It gathers the physical telemetry to assist human users or feed into autonomous systems.



Technology investment stack tier:
Application tier




Application tier

This segment serves as the intelligence demand interface where human intent and corporate workflows are aggregated and executed. It represents the transition from seat-based SaaS licenses to autonomous digital labor executing horizontal and vertical tasks across the enterprise. This tier acts as the ultimate top-of-funnel for requested digital labor, generating the massive compute requirements that cascade down the rest of the architecture.

- **Layer 22—consumer AI:** The aggregate human-demand layer. This software tier encompasses B2C applications, social networks, and entertainment platforms where global human attention is monetized. It functions as the top-of-funnel for the digital consumer economy, aggregating intent and driving the massive compute requirements cascading down the stack as individuals’ wants and needs are addressed.
- **Layer 21—industrial AI:** The proprietary industry logic base. It includes highly specialized OS workflows and platforms tailored to complex and regulated sectors like healthcare, finance, defense, and heavy manufacturing. It operates as the proprietary intelligence layer for end-market AI and houses nonreplicable institutional data that generalized models cannot access due to regulatory and privacy constraints.
- **Layer 20—enterprise AI:** The horizontal corporate workflow engine. This layer categorizes the transition of legacy SaaS into digital labor across two fundamental macroeconomic functions: revenue optimization (growth, service, knowledge management, and creative workflows) and cost optimization (operations, asset management, human capital, supply chains, finance, analytics, and governance). Investment is aggressively rotating out of static user licenses and into autonomous agents that independently execute these horizontal enterprise tasks, creating fundamentally original opportunities in corporate margins and operating leverage.



Technology investment stack tier:
Intelligence tier




Intelligence tier

This segment houses the core intelligence models, foundational data architectures, and operational guardrails of the AI economy. It is where raw data, both structured and unstructured, is governed, processed, and transformed into autonomous, agentic reasoning. Enterprises must deploy heavy capital here to ensure their digital workforce is accurate, compliant, and consistent.

- **Layer 19—agentic AI:** The executive reasoning layer. This consists of software frameworks executing multistep goals autonomously. It serves as the cognitive middle management of the stack. This layer also heavily encompasses AgentOps for operational management and AI TRiSM for fundamental model governance of execution.
- **Layer 18—generative AI:** The core cognitive engine. It is built upon foundation models capable of translating massive unstructured datasets into net-new outputs, including logic, text, audio, visual, or code. To securely operationalize this intelligence, enterprises allocate significant software budgets to MLOps and LLMOps at this layer.
- **Layer 17—data fabric:** The enterprise memory bank. Encompassing relational databases, data lakes, and high-dimensional vector databases, it serves as the critical retrieval system that grounds AI models. This layer is strictly governed by data governance and MDM tools to protect corporate IP and enforce access controls.



Technology investment stack tier:
Infrastructure platforms




Infrastructure platforms

This segment functions as the overarching control plane and delivery network for the AI factory. It is responsible for orchestrating, virtualizing, and securing the flow of compute across distributed environments. Without extensive development and redundancies in these critical routing and abstraction layers, the elastic deployment of complex AI workloads would collapse.

- **Layer 16—orchestration:** The traffic conduction layer. Driven by container management systems, API gateways, and CloudOps/observability platforms, it autonomously scales and balances workloads. This layer also physically houses the automated execution pipelines for DevOps and DevSecOps.
- **Layer 15—virtualization:** The foundational software abstraction layer. Utilizing hypervisors and virtualized operating systems, this layer is responsible for slicing bare-metal physical servers into highly efficient and isolated virtual machines, ensuring enterprises squeeze maximum utilization out of scarce GPU clusters.
- **Layer 14—secure delivery:** The zero-trust application delivery and cybersecurity perimeter. This layer guarantees that compute traffic is securely routed from the datacenter core to the intelligent edge. By converging global load balancing, critical API security, and continuous network telemetry, it acts as the defense architecture shielding the AI factory from automated botnets and multi-cloud vulnerabilities.



Technology investment stack tier:
Datacenter tier




Datacenter tier

This segment represents the physical housing, thermal regulation, energy delivery, and hyperscale compute machinery powering the AI economy. It is a highly capital-intensive tier constrained by asset scarcity, including real estate footprints, power generation bottlenecks, basic component availability (GPUs, RAM, storage, and more). Owning hard assets in this segment provides a structural moat in the intermediate term while these shortages persist. This is also one of the areas where significant political opposition has begun to coalesce. This stands to both increase risk at development and increase the value of assets in production.

- **Layer 13—hyperscale cloud service:** The centralized computing tier. This is where hyperscalers and specialized cloud providers host multi-tenant compute and platform services. It serves as the fundamental processing engine of the global economy, capturing the majority of AI software spend by monetizing intelligence. Revenue here is aggressively driven by inference as a service and high-volume token generation, acting as the computational foundry required to serve foundation models at scale. This is currently dominated by the hyperscalers Amazon, Microsoft, and Google, as well as additional participants, including Meta and CoreWeave.
- **Layer 12—colocation & real estate:** The physical datacenter shell. This layer is provided by landlords operating secured wholesale and colocation facilities. It serves as the literal real estate underpinning the digital economy, heavily constrained by the availability of power and water infrastructure, as well as by localized resource agreements for both.
- **Layer 11—IT hardware:** The physical computing machinery. This layer includes enterprise servers, storage arrays, networking switches, and high-density GPU chassis. It provides the tangible silicon housing and network backplanes required to physically process the trillions of parameters within modern AI models.
- **Layer 10—power & thermal equipment:** The vital datacenter regulation layer. It encompasses uninterruptible power supplies, precision cooling, and advanced liquid-cooling technologies. It prevents catastrophic hardware failure by



Technology investment stack tier:
Semiconductor



regulating the immense heat and electrical draw generated by tightly packed GPU clusters, enabling optimal use of these massively expensive assets and allowing them to deliver uninterrupted digital labor to the layers above.


Semiconductor tier

This segment is the silicon-based foundation of all digital computation. It captures the deeply consolidated supply chains, fabless design monopolies, and bleeding-edge physics required to produce intelligent digital labor. Value capture is protected by nearly insurmountable capital investments, deep tech requirements, specialized technical expertise, and geopolitical strategic importance.

- **Layer 9—quantum technologies:** The superposition processing layer. This layer leverages quantum mechanics, superconducting qubits, and trapped ions for probabilistic computation. In its early stages, this will act as a strategic sovereign hedge designed to defend from quantum computing-based cyberattacks and solve specific high-theoretical math/science and material science problems that classical supercomputers cannot. This technology is still under development and far from proven. When it is widely deployed, many of our assumptions about the AI economy and ecosystem will radically change.
- **Layer 8—chip design:** The silicon processing layer. It includes fabless and IDM chip manufacturers. These companies design instruction set microarchitectures for CPUs, GPUs, and custom ASICs. This is the highest-margin segment of the physical stack, dominated by monopolies that possess proprietary software layers in their design architectures.
- **Layer 7—semiconductor supply chain:** This is the semiconductor fabrication production base, which consists of multibillion-dollar foundries, lithography, materials, intellectual property, wafer fabrication equipment, advanced packaging, and testing facilities. This layer operates as a highly integrated global supply chain with limited fail-safe redundancy. As a result, it is a highly visible geopolitical chokepoint responsible for the uneven development of AI computing globally.



Technology investment stack tier:
Physical infrastructure tier




Physical infrastructure tier

This segment provides the macro-scale utilities and physical distribution lines required to sustain the technology sector. It encompasses the global telecom grids, raw baseload power generation, and component assembly lines that serve as the ultimate bottleneck for AI scaling. Capital must flow here first to clear the constraints hindering the deployment of upper-layer intelligence.

- **Layer 6—manufacturing infrastructure:** The hardware manufacturing grid. This layer comprises electronics manufacturing services, original design manufacturers, and industrial tooling for both IT and operational technology. It acts as the macro assembly line and component distribution network, enabling original equipment manufacturers across all sectors to seamlessly outsource or scale physical production for everything from hyperscale compute chassis to autonomous robotics. The same robotics produced by this layer will further advance its capabilities in a virtuous cycle.
- **Layer 5—communication infrastructure:** The data communication grid. It consists of physical cell towers, subsea transoceanic cables, terrestrial fiber-optic networks, and low-Earth orbit communications satellites. It is the essential physical highway connecting the centralized intelligence of the hyperscale cloud to the distributed nodes of the intelligent edge.
- **Layer 4—power infrastructure:** The new energy grid. This layer encompasses baseload power generation facilities, high-voltage transmission lines, and utility-scale energy storage that is now scaling to address massive new demands. It produces the necessary electricity, serving as the ultimate limiting factor for the entire digital and cognitive economy above it.



Technology investment stack tier:
Services tier



Services tier

This segment captures the elite human capital and physical labor required to integrate and construct the AI economy. It ranges from strategic management consulting to the pouring of concrete for new datacenter facilities. This tier reinvigorates labor opportunities in complex engineering, physical construction, and enterprise change management.

- **Layer 3—IT services:** The digital plumbing and hardware deployment layer. It comprises value-added resellers and managed service providers. It provides the specialized project labor required to unbox, rack, network, and configure the physical servers and execute complex cloud migrations.
- **Layer 2—engineering & construction:** Heavy physical construction. This layer deploys human labor, heavy equipment, and logistics to execute macro-scale infrastructure projects. It is the organizational force responsible for physically constructing hyperscale datacenters, semiconductor fabs, and power plants that house the data revolution.
- **Layer 1—professional services:** The strategic layer of humans providing advisory and systems integration services. It consists of elite management consultants deploying human intellect to orchestrate enterprise-wide digital transformations, re-engineer corporate workflows, and manage organizational change.



Investment risks

The transition to digital labor carries profound structural and execution risks across all eight segments of the stack. Investors should assess these vulnerabilities, as failure at any single layer can cascade through the entire digital-to-physical value chain.

Global macro effects

While the stack provides a structural framework for mapping the AI economy, its development is also being shaped by external forces that influence how capital is deployed, how systems are built, and where adoption occurs. Sovereign capital is shaping this development and how it scales across regions. Governments are no longer only setting regulatory frameworks; they are active buyers and financiers of AI infrastructure from semiconductors to datacenters and national cloud platforms. This introduces a new durable, noncyclical demand driver, where capital deployment is driven not only by productivity but also by national security priorities, industrial policy, and technological self-sufficiency.

The stack is gradually splitting into geopolitically aligned ecosystems, anchored by a dominant US-centric stack and a rapidly developing Chinese alternative. The US maintains leadership in LLMs and the accelerator ecosystem, but export controls, capital restrictions, and supply chain localization are accelerating domestic substitution in China. Semiconductor Manufacturing International Corporation (SMIC), HiSilicon, and Huawei are excellent examples of how US trade restrictions have accelerated Chinese domestic chip innovation and increased long-term price competition for US-aligned manufacturers. In this scenario, capital investments are determined by political alignments rather than by market forces. Likewise, supply chain capacity is artificially expanded based on the mutually exclusive geographic logistics of two unaligned groups rather than optimized on supply chain continuity and redundancy for the entire system.

Endpoint liability and software cannibalization

At the top of the stack, the physical deployment of intelligent edge-like robots, autonomous vehicles, drones, and industrial machines introduces unprecedented liability and physical safety risks, alongside expanding the attack surface for bad actors targeting unsecured edge networks. The capital required to insure, fix, and physically recall



compromised fleets will severely compress operating margins for early commercial adopters. Moreover, regulatory frameworks governing autonomous machine behavior remain largely undefined. Widespread industrial deployment is limited until clear liability standards are established.

Within the application tier, the transition from seat-based SaaS to outcome-based digital labor will trigger deflationary pricing pressure, aggressively cannibalizing legacy recurring revenue models across horizontal and vertical software markets. Incumbents who fail to transition their billing structures will face revenue cliffs as enterprise buyers consolidate and cancel redundant software licenses. Consequently, a massive wave of vendor consolidation is inevitable as legacy software companies are forced to acquire agentic-native startups simply to survive the transition.

Regulatory headwinds and infrastructure vulnerability

Beneath the application layer, foundation models and agentic workflows face severe regulatory headwinds regarding data provenance, copyright infringement, and the enterprise liability associated with autonomous algorithmic hallucinations. Strict data sovereignty mandates, such as China's Data Security Law (2021) and India's DPDP Act (2023), now restrict cross-border data flows, while pending sovereign legislation like the EU Data Act (2023/2024) threatens to retroactively invalidate proprietary training sets that incorporate protected industrial data. Furthermore, transparency mandates like the federal CLEAR Act (2026) and California's AB 2013 (2024/2026) require detailed public reporting of training sources, creating a provenance audit that could force costly, region-specific model retraining and severely stall global commercial deployment.

The computing power needed to generate digital labor intelligence is controlled by the infrastructure software tier, where orchestration, virtualization, and secure delivery each create single points of failure, as demonstrated in recent outages by hyperscalers affecting thousands of businesses globally. Specifically, localized virtualization or routing



errors can easily cascade into global network outages. The economic fallout of these concentrated outages will drive enterprise pushback, forcing SaaS vendors to build multi-cloud redundancy architectures. The mandate for resiliency will add to infrastructure operating costs, favoring networking incumbents who possess global distribution and maintenance footprints.

Power constraints and silicon supply chain fragility

At the core of the physical compute engine, infrastructure expansion is structurally capped by hard constraints, specifically the lack of available baseload power, localized thermal limits, and aggressive zoning pushback against multi-gigawatt builds. This severe supply-demand mismatch threatens to throttle the compute growth of incumbent hyperscalers, emerging neoclouds, and dedicated high performance computing AI factories as they fiercely compete for the same scarce energy agreements. Consequently, to maintain operating leverage and higher computing capacity per rack (compute density), capital allocation must aggressively pivot toward bespoke, capital-intensive infrastructure. This includes on-site power generation and battery storage systems, as well as advanced liquid-cooling architectures, which all substantially elevate execution risk.

Supporting these datacenters is a semiconductor layer that remains critically vulnerable to geopolitical risks. Extreme capital requirements and geographically concentrated fabrication plants create fragile, single-node supply chains that are instantly disrupted by rapid regional destabilization, halting global IT hardware production. Moreover, the physical limitations of Moore's Law currently force semiconductor design into highly experimental multi-die packaging, significantly increasing the probability of cumulative manufacturing defects and stranding software investments dependent on next-generation compute.



Macro-engineering delays and human capital deficits

Finally, advanced computing AI datacenters face massive, real-world execution hurdles when trying to be built. The engineering required to sustain this expansion faces severe execution delays due to multiyear utility permitting processes and global shortages in both raw materials like copper and finished goods like electrical transformers. These timeline extensions will structurally inflate baseline capital expenditure requirements, delaying the break-even models for mega-project datacenter infrastructure funds. Unless sovereign governments actively streamline environmental regulations and subsidize heavy manufacturing, the entire cognitive expansion risks stalling out against the rigid limits of physics and material availability.

Compounding these material limits is a strict bottleneck in human capital within the services segment. There is a severe lack of specialized EPC labor required to integrate these complex environments. Escalating wage premiums for these scarce technical specialists will continuously erode project returns and throttle the deployment pace of the broader digital economy. This structural talent shortage will ultimately force hyperscalers and engineering firms to poach each other's workforces, driving a cycle of human capital inflation, as observed in the high-end AI model development labor force.



Brave new world

As discussed in [SaaS Is Dead, Long Live SaaS](#), traditional software valuations will continue to face pressure. Foundation models will continue to commoditize both application development and knowledge-based professional verticals. However, software companies that can provide a comprehensive digital labor force to automate entire employee workflows with outcome-based pricing models will likely become the new incumbents. In [Through the Looking Glass: The Race to Build Enterprise AI](#), we introduced new frameworks to help stakeholders evaluate startup and incumbent software companies' readiness for the transition to digital labor. We also stressed to IT managers the importance of enterprises building agent operations to close the autonomy gap when using these services.

In this note, we combine advanced computing infrastructure and SaaS into a comprehensive technology stack. Each layer represents a significant investment and technological foundation for the digital labor economy. This stack maps where AI operates and how capital will flow from upstream autonomy to downstream infrastructure once the AI economy is fully achieved.

We see the transition to digital labor across corporate enterprises as the next significant capital expenditure super-cycle. This shift is misconstrued as the end of nonfarm payroll employment and the vast revenue on which it is based. Recent market turbulence has caused almost all software assets to be re-rated down from historical ranges. This incentivizes vendors to improve both product-market fit and pricing models to better align with the future economy. For incumbent software leaders, this transition is not without peril. Shifting from predictable seat-based ARR to potentially more volatile, outcome-based consumption models is causing investor concern. We expect short-term margin guidance to focus on agentic revenue growth. Overall performance, however, may be muted as vendors assess legacy product, pricing, and investment plans. Only those incumbents with the leadership and talent to aggressively fund native digital labor architectures while achieving early success in new product growth will maintain account control and investor interest in the near term.



Going forward, corporate and industry defensibility will require owning and managing the massive physical hardware infrastructure at the bottom of the stack (layers 4 to 13) or owning the trusted customer data and context-specific business logic at the top (layers 20 to 25) that delivers digital labor.

Ultimately, digital labor does not eliminate the need for human capital—it accelerates it. Killer apps have always marked the inflection point for the mass adoption of every technology platform. With digital labor, the killer app is human, and this inflection point is now.

Soon, trade professionals, small business owners, enterprise employees, and industrial engineers will be equipped with teams of autonomous agents to optimize productivity and income growth. AI fundamentally transforms global industries by re-engineering companies, industries, and economies at the human level.

During this time, innate human traits, such as creativity, empathy, flexibility, and trust building, will become exceptionally valuable. When agents optimize cost, speed, data processing, and physical logistics, the only remaining variables that can drive market differentiation are human. The highest valued assets in this economy will be individuals who can leverage agents to establish productive customer relationships and find novel growth opportunities.

At the highest level, investors must fundamentally align their investment theses on founder-led product-market fit. As expert-level coding becomes instantaneous, only human understanding of human challenges will drive future technology innovation. Consumers and enterprises will pay a premium for human provenance and the assurance of human accountability. IaPs will offer WaaS priced by safety, reliability, execution, and peace of mind. They will use their digital fleets to reduce costs, maintain continuous uptime, and address their clients' emotional concerns. Finally, they will deliver hyper-personalized experiences that AI can never orchestrate alone.



In the long term, individuals and enterprises with the highest operating margins and growth will be those with the greatest understanding of end-user pain points. We see successful investments over the next decade focused on platforms that enable comprehensive SaaS digital labor, empower human orchestration of agents, and enable human-to-human trust building with guaranteed, repeatable outcomes and new product innovation. This is the foundation of the next great expansionary super-cycle. To capture this growth, stakeholders must actively pursue investments across all levels of the stack that enable digital labor, recognizing that while intelligence scales nearly instantaneously with compute, the super-cycle remains governed by physical scarcity limits. Companies led by management that do not anticipate a future WaaS economy will be replaced by those that do.




References

- 1: ["Forecast Analysis: Enterprise IT Spending Across Vertical Industries, Worldwide," Gartner, October 21, 2025.](#)
- 2: Ibid.



PitchBook provides actionable insights across the global capital markets.

Additional research:



Q1 2026 Analyst Note: SaaS Is Dead, Long Live SaaS

Download the report [here](#)



Q1 2026 Analyst Note: Through the Looking Glass: The Race to Build Enterprise AI

Download the report [here](#)

[PitchBook Insights](#) is an online compendium of in-depth data, news, analysis, and perspectives that shape the private capital markets.

PitchBook subscribers enjoy exclusive access to a comprehensive suite of private market insights, including proprietary research, news, data, tools, and more on the [PitchBook Platform](#).

COPYRIGHT © 2026 by PitchBook Data, Inc. All rights reserved. No part of this publication may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, and information storage and retrieval systems—without the express written permission of PitchBook Data, Inc. Contents are based on information from sources believed to be reliable, but accuracy and completeness cannot be guaranteed. Nothing herein should be construed as any past, current or future recommendation to buy or sell any security or an offer to sell, or a solicitation of an offer to buy any security. This material does not purport to contain all of the information that a prospective investor may wish to consider and is not to be relied upon as such or used in substitution for the exercise of independent judgment.

Nizar Tarhuni

Executive Vice President of Research and Market Intelligence

Paul Condra

Senior Director, Global Head of Private Markets Research

James Ulan

Director, Industry & Technology Research

Report created by:

Rudy Torrijos

Director, Industry Research

Derek Hernandez

Senior Research Analyst, Enterprise SaaS and Infrastructure SaaS

Dimitri Zabelin

Senior Research Analyst, Artificial Intelligence and Cybersecurity

Oscar Allaway

Senior Data Analyst

Jenna O'Malley

Senior Graphic Designer

Learn more about [PitchBook's Institutional Research team](#).

Click [here](#) for PitchBook's report methodologies.