


Launch Report: Advanced Computing

PE and VC trends and investment strategies

Q3
2025





Contents

Executive summary	3
Introduction	6
Market scenario	8
Market risks	9
Market growth areas	10
Recommendations to stakeholders	12
Overview and scope	14
Total addressable market	18
Growth drivers	20
Pivots and catalysts	24
Leadership matrix	27
Segment overview	34
Semiconductor supply chain	34
Semiconductor chip design	35
Quantum computing	35
Datacenter infrastructure	38
Deal and exit activity	42
Advanced computing PE and VC ecosystem market map	61
Advanced computing PE investor map	62
Advanced computing VC investor map	63
Taxonomy	64
Glossary	70
Methodology	82

Institutional Research Group

Analysis



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

Data

Oscar Allaway Data Analyst

pbinstitutionalresearch@pitchbook.com

Publishing

Report designed by **Jenna O’Malley** and **Megan Woodard**

Published on December 10, 2025




Executive summary

This report launches our coverage of the advanced computing vertical. Advanced computing comprises all end markets that enable artificial intelligence (AI) to scale. This includes both semiconductor and datacenter markets across both operational and information technology (IT) vendors. Advanced computing is one of two verticals within PitchBook's advanced technologies research domain. Coverage of PE software will be launched in 2026.

We see hyperscaler capital expenditures as the single defining metric for investors and vendors to gauge the growth and sustainability of the advanced computing ecosystem. We believe the current expected hyperscaler capital expenditure level for 2025 and 2026 is realistic and indicative of conservative strategic-planning expectations. We view power generation as the primary gating factor to revenue growth for advanced computing vendors through 2030. Macroeconomic and semiconductor cyclicalities are unlikely to affect capital expenditure deployment as reported end-customer return on investment (ROI) remains significant.

Current capital expenditures are primarily focused on increasing the fidelity of generative foundation models

Hyperscaler capital expenditures (\$B) by quarter



Sources: Company filings and PitchBook • Geography: Global • As of November 26, 2025

Note: Data is based on PitchBook analyst calculations. The Big Five hyperscalers are Microsoft, Amazon, Google, Meta, and Oracle.



EXECUTIVE SUMMARY

(training and inference). Without this achievement, end-user productivity from generative AI (GenAI) models will be meaningfully limited beyond 2027. Agentic AI and physical AI development will be significantly curtailed as well.

The primary beneficiaries of hyperscaler investments within the advanced computing sector are graphics processors, network processors, server microprocessors, AI accelerators, leading-edge-node semiconductor supply chain vendors, power-generation equipment, thermal management equipment, and neoclouds. We do not see quantum computing as a commercially viable technology through 2030.

Key takeaways

To assist investors and stakeholders in navigating this transformational shift, we offer these key takeaways that will shape the advanced computing market over the next five years:

- **Capital expenditure magnitude:** Hyperscaler capital expenditures are projected to reach a cumulative \$6.4 trillion through 2030, representing the single largest deployment of investment capital in human history.

- **Elasticity of intelligence:** Demand for processors, projected to generate \$2.8 trillion in cumulative revenue, is effectively uncapped. Intelligence scales with computing power and is limited only by problem complexity.
- **Architectural efficiency:** The industry-wide transition to Mixture-of-experts (MoE) architectures has drastically lowered inference costs, making “expert-level” intelligence economically viable and triggering the current infrastructure boom.
- **Fidelity risk:** Language model fidelity (LMF) is the single most critical performance metric, as a failure to achieve perfect accuracy will relegate AI to creative tasks and stall the high-value agentic and physical AI markets.
- **The 60% threshold:** Surpassing the 60% score on the ARC-AGI-2 benchmark (average human intelligence) is the critical technical milestone required to validate the labor-substitution thesis and justify the massive capital outlay.
- **Semiconductor supercycle:** The semiconductor sector has entered a secular supercycle driving \$3.7 trillion

in cumulative datacenter chip revenue through 2030, effectively decoupling the industry from traditional cyclical downturns.

- **Power constraints:** Power availability has replaced silicon as the primary constraint on growth, driving \$1.2 trillion in cumulative spending on power and cooling infrastructure through 2030 to address three-to-five-year grid interconnection queues.
- **Thermal obsolescence:** Legacy air-cooled datacenters (limited to around 30 kW per rack) are now technologically obsolete for AI clusters, necessitating a massive retrofitting cycle to liquid-cooling technologies to prevent thermal throttling.
- **Supply chain fragility:** Geopolitical concentration remains an existential risk, as the entire advanced computing stack depends on TSMC’s capacity, with supply chain sovereignty requiring at least a decade to achieve.



EXECUTIVE SUMMARY

- **Not dot-com:** Unlike the speculative dot-com era, current hyperscaler capital expenditures are grounded in broad-based customer demand and backed by record operating cash flows, mitigating fears of an imminent asset bubble.
- **The rise of neoclouds:** Specialized “neoclouds” have emerged as a critical market segment, capitalizing on the “time-to-power” gap to offer bare-metal GPU availability faster than generalist hyperscalers.
- **The importance of 2 nm:** The transition to 2-nm process nodes is a physical requirement to meet power-efficiency targets, disproportionately benefiting leading-edge foundries within a \$3.1 trillion supply chain market.
- **PE strategy:** Value creation for PE is shifting from passive datacenter real estate to the active consolidation of critical “industrial” supply chain vendors in liquid cooling and power distribution.
- **VC pivot:** Alpha generation has migrated from general foundation models to the “reasoning layer” and physical AI, where proprietary industrial data creates defensive moats against commodity chatbots.
- **Enterprise mandate:** Enterprises should abandon building foundation models in favor of fine-tuning and retrieval-augmented generation (RAG) architectures to dramatically reduce the total cost of ownership while treating computing supply as a scarce physical commodity requiring immediate multiyear reservation contracts.



Introduction

OpenAI launched ChatGPT-3.5 on November 30, 2022, exactly three years before the writing of this report. On November 22, just one week earlier, NVIDIA reported datacenter revenues of \$3.8 billion. Today, NVIDIA reports quarterly datacenter revenues of \$51 billion, a 13.4x increase in three years.

Technologists, macroeconomists, the world's largest corporations, sole proprietorships, sovereigns, and the general public universally agree that AI has changed everything. In Oracle's Q2 2025 earnings conference call, Chief Technology Officer Larry Ellison said, "Not everyone fully grasps the extent of the tsunami that is approaching." We believe he is correct, and in fact might be understating the radical transformation we will see within labor-force skill sets and the investment needed

to support it. The dinosaurs would have survived a mere tsunami—this is much bigger.

The GenAI era began on January 20, 2025, with the introduction of DeepSeek's R1. DeepSeek proved that using an MoE architecture for language models can dramatically improve customer fidelity and performance per dollar. This was the proverbial asteroid that changed everything. Today, we are seeing unprecedented capital investments from the world's largest, most profitable companies, because unlike in 2022, 2023, and 2024, we now have a road map to perfect LMF. Service providers, enterprises, and operational technology (OT) companies can now invest in and plan for AI that works as intended.


The goal of this investment is not superintelligence but return on capital, margin expansion, and revenue growth. Capital investments will continue to approach hyperscaler operating cash flow until AI approaches expert level. We do not see this happening until 2030.

We launch our coverage of advanced computing with this report. In the following sections, we provide an overview of the vertical and its market segments. Going forward, we will report on deals, exits, valuations, and market trends. This launch report provides investors and stakeholders with the framework to track and analyze advanced computing growth and risk during the greatest societal change since the assembly line and electricity distribution.



INTRODUCTION

Big Five hyperscaler operating cash flow and capital expenditures



Sources: Company filings and PitchBook • Geography: Global • As of November 26, 2025
Note: Data is based on PitchBook analyst calculations. Hyperscaler operating cash flow is derived from all company revenues, not just cloud services. The Big Five hyperscalers are Microsoft, Amazon, Google, Meta, and Oracle.



Market scenario

Semiconductors

- 2025 marked the beginning of an unprecedented five-year semiconductor supercycle. Inventory correction in the industrial and automotive chip sectors simultaneously peaked with a broad-based acceleration in demand for processors and accelerators for datacenters.
- We calculate that the 2030 semiconductor total addressable market (TAM) has increased by \$600 billion in the past 18 months and that cumulative hyperscaler capital expenditure forecasts will rise to \$6.4 trillion over the next five years.
- Datacenter semiconductors will represent 63% of the total semiconductor market in 2030, up from 11% in 2022. AI accelerators, server processors, and embedded processors are the primary beneficiaries of growth in GenAI, agentic AI, and physical AI applications, respectively.
- Leading-edge foundries and advanced wafer fabrication equipment (WFE) companies will disproportionately benefit from the datacenter semiconductor boom. Foundries with 2-nm node capacity and extreme ultraviolet (EUV) equipment are needed to address the vast power efficiency and processing requirements of AI datacenters.

Datacenters

- The introduction of DeepSeek's R1 language model based on an MoE architecture in early 2025 was a groundbreaking milestone in AI computing. It marked the "big bang" event in hyperscaler IT capital expenditures, which we now forecast to approach \$4.7 trillion over the next five years.
- Power and thermal management (cooling) equipment is the gating factor to datacenter investment deployment. We forecast that power and cooling will reach \$1.2 trillion in cumulative spending over the next five years. Currently, we estimate that power and cooling account for 17% of the cost to build a new datacenter. This is expected to grow to 22% as higher-performance chips are introduced in 2027 and 2028.
- Hyperscaler capital expenditures are not slowing, nor are they akin to the dot-com bubble. Capital expenditures are currently at 60% of the free cash flow of hyperscaler companies. This is a key metric to watch. The largest companies in the world are hyperscalers, and they are making strategic-planning decisions based on broad customer demand. During the dot-com era, startups using VC and public capital generated their own circular new economy without broad-based technology investment.

- The 2025-2030 period will mark the commercialization of the GenAI era. Superintelligence or even artificial general intelligence (AGI) is not the goal of hyperscalers or foundation model companies; return on invested capital is. More powerful graphics processing and AI accelerators, coupled with advancements in MoE architectures, are enabling larger-scale training and increased accuracy of inference. We expect nearly \$2.1 trillion to be spent on processors over this period to accomplish this goal.
- LMF, defined as the accuracy of a language model in following a user's intended instructions, is a key driver for further growth. Recent models have shown dramatic increases in expert intelligence, improving customer utility. We view LMF and "tokens per dollar per watt" as the primary key performance metrics for hyperscalers, end users, and investors.
- Both agentic AI and physical AI are highly dependent on increasing LMF. Without perfect fidelity, both agents and robots may randomly misunderstand their tasks. As fidelity is achieved, we see enterprises (through agentic AI) and manufacturing (through physical AI) accelerating capital investment further from 2030 through 2035.



Market risks

We expect that the elasticity of intelligence will drive a projected \$6.4 trillion in cumulative datacenter spending through 2030. There are four significant risk factors that could slow growth and dislocate capital. However, each could be a growth driver if current limitations are ameliorated.

Power and cooling

- **Grid capacity:** Power availability is the primary gating factor for AI scaling through 2030. US and European grids face interconnection queues of three to five years. This caps the total number of GPUs that can be deployed.
- **Sustainability versus reality:** Hyperscalers pursuing net-zero carbon targets alongside gigawatt-scale power capacity may adjust their growth expectations. Reliance on natural gas or diesel backup creates regulatory risk. At scale, datacenters are unregulated fossil-fuel power plants in the eyes of environmental agencies.
- **Thermal limits:** Legacy air cooling (limited to around 30 kW per rack) is obsolete for modern Blackwell-era clusters (more than 100 kW per rack). Failure to secure liquid-cooling systems risks thermal throttling and hardware failure. Local

governments are increasingly blocking datacenter permits due to water usage concerns, capping the number of GPUs that can be deployed.

Language model fidelity

- **Stranded assets:** Failure to achieve higher-than-average adult human intelligence to solve unfamiliar problems (fluid intelligence) with higher computing resources will reduce the need for such higher computing resources. Hyperscalers will deploy capital more slowly as the usefulness of GenAI will be limited to knowledge-based (crystallized intelligence) industries such as healthcare, life sciences, drug discovery, fintech, and software engineering.
- **Human overseers:** Hallucinations or misunderstandings of user intent force enterprises to maintain their AI verification processes. If companies cannot remove human oversight, margin expansion and productivity will fail to materialize.
- **Safety first:** Autonomous machines may misinterpret safety-critical instructions. High-profile accidents caused by low-fidelity agents acting physically in the real world could lead to strict regulatory crackdowns, slowing down investments in industrial AI.

Supply chain sovereignty

- **The Taiwan problem:** TSMC dominates leading-edge-foundry capacity. Any kinetic or blockade event in the Taiwan Strait would halt the production of 90% of AI accelerators and logic integrated circuits (ICs).
- **Export-control retaliation:** Continued US restrictions on H100/H200 exports may accelerate Chinese domestic innovation (such as with Huawei and SMIC), potentially bifurcating the global standard for AI hardware and software stacks.

Concentration

- **Market concentration:** 80% to 85% of all advanced computing spending will come from five companies: Amazon Web Services (AWS), Microsoft, Google, Meta, and Oracle. A strategic pivot by just one hyperscaler would materially impact the entire semiconductor value chain. Hyperscaler capital expenditures among the Big Five are currently at a record 77% of operating cash flow. If AI-native vertical software systems and enterprise application software do not deliver increasing revenue growth and profitability by 2027, shareholders may force a capital expenditure pullback from any one of these companies.



Market growth areas

Investors

Venture capital

The “reasoning layer” is the next focus in innovation. Agents will eventually replace specific human skill sets in highly specialized knowledge-based verticals. Physical AI robotics and industrial automation using proprietary data will create significant growth opportunities for industrial companies. Neoclouds offer immediate time to power and GPU availability that hyperscalers struggle to provision for midmarket clients.

Private equity

Datacenters are long-lived AI factories, not just giant warehouses of IT hardware. PE investors should shift their focus from real estate (shell and power) to AI-critical supply chain suppliers. Fragmented vendors in liquid cooling (coolant distribution units and manifolds) and power distribution (switchgear and transformers) will need to be consolidated. Power and cooling supply chain vendors will benefit from economies of scale, proving faster time to power than point solution vendors. The moat is no longer the building shell, but the ability to handle densities over 50 kW, secure gigawatt power feeds, and deliver thermal management solutions.

Public equity

Hyperscalers, power and cooling equipment, AI processors, foundries, leading-edge semiconductor equipment, and electronic design automation (EDA)/intellectual property (IP) remain key beneficiaries. The current GenAI era will see significant investments over the next five years. Rapid deceleration in the upward revision of hyperscaler capital expenditures and stagnant chip-company margin expansion will be possible indications of slowing demand from end customers or waning productivity gains. If expert intelligence revenue lags capital expenditures, expect a short-term pullback in hardware spending.

Service providers

Investment banks

Investment banks should expect consolidation in the advanced-packaging supply chain. As organic substrates hit physical limits, manufacturers of glass core substrates and silicon photonics (optical interconnects) are prime targets for acquisition by foundries and outsourced semiconductor assembly & test companies (OSATs) seeking to break the memory wall.

Smaller OSATs and specialized power and cooling equipment companies will be consolidated by major equipment suppliers seeking end-to-end capability.

Investment banks should look to structure joint ventures between US tech giants and sovereign wealth funds (such as those in the United Arab Emirates, France, and Japan) to build localized sovereign clouds. These deals require navigating complex government policies in nondomestic investment and export control.

Commercial banks/private credit

Datacenters are seeking independence from the public grid. AI datacenters will require massive on-site power generation. Power infrastructure now constitutes a significant portion of an AI facility’s capital expenditures, requiring specialized financing for large-scale on-site gas turbines and fuel-cell installation.

The datacenter asset-backed securities market can now fund large-scale AI projects. As hyperscalers sign long-term leases for AI clusters, banks can package these predictable cash flows into investment-grade debt products for institutional investors.



MARKET GROWTH AREAS

Technology vendors

Sovereign AI cloud capital expenditures are forecast to grow faster than the overall hyperscaler market. Nations across the globe are mandating domestic data residency and require turnkey AI solutions. Technology vendors should partner with local telecommunication companies and governments to build air-gapped, domestic AI clouds that satisfy strict data sovereignty laws.

Corporate strategy

Hybrid computing optimizes costs. Enterprises should pursue reserved instances for their steady-state baseload inference and use spot instances or neoclouds for burst training. This hybrid approach can dramatically reduce total computing costs.

Enterprises should not build their own foundation models unless they have proprietary data that constitutes a defensive moat. Instead, they should leverage open weights or proprietary APIs and focus engineering on fine-tuning and RAG for internal data. Using these architectures for internal workflows can dramatically reduce the total cost of ownership for inference compared with relying exclusively on frontier models for every query.

Enterprises should secure long-term supply chain contracts. Current expert-level intelligence (crystallized intelligence) is sufficient for high-ROI automation in coding, customer support, and most research & development (R&D). Enterprises should treat computing supply like a physical commodity and sign multiyear capacity reservation contracts now, as the backlog for Blackwell-class silicon will persist through 2026.



Recommendations to stakeholders

Capital investment in advanced technology industries will be unprecedented in scope. Commercial banks, private credit, venture capital, corporate venture capital, hyperscaler free cash flow, and public capital markets will fund or finance most of this investment. We recommend that stakeholders consider focus areas across the entire advanced technology stack and avoid industry concentration risk.

Advanced computing, as covered in this report, and advanced applications (AI and software as a service) make up the entire advanced technology stack. With this framework, we identify direct and indirect competitors and research end-market performance trends. We recommend that stakeholders

analyze quarterly changes, public company commentary, and private company deal and exit flows to assess end-market demand risk.

Within our advanced computing research, we track deal financing, exits, valuations, and product trends across 40 categories. With nearly \$1 trillion in capital investment by hyperscalers per year starting in 2027, each of these end markets will see accelerating revenue, product innovation, and financings. We recommend stakeholders analyze quarterly valuations, deal flow, and management changes to assess risk to product-market fit.

Companies with the best technologies are seldom long-term leaders in their industries. Companies with management focused on building strategic customer partnerships, reducing product time to value, and expanding profitability are market-share winners. We recommend stakeholders analyze quarterly staff adjustments, head count growth, operating models, and pricing trends to assess management execution risk.



RECOMMENDATIONS TO STAKEHOLDERS

Advanced technology stack and leading providers

Advanced applications	Vertical software systems: Fintech, healthcare, industrial, aerospace & defense, transportation, manufacturing, retail, agriculture	Software
	Application software: Microsoft, Oracle SAP, ServiceNow, Salesforce, HubSpot, Adobe, ADP, Intuit, Workday, Atlassian	
	Infrastructure platforms: Microsoft, Oracle, IBM, Cisco, Broadcom, Palo Alto Networks, ServiceNow, CrowdStrike, Fortinet	
Advanced computing	Physical AI systems: Siemens, Schneider Electric, Honeywell, Rockwell Automation, FANUC, Yaskawa, KUKA, ABB, Tesla, Waymo	Intelligence
	Agentic AI systems: Microsoft, Google, AWS, NVIDIA, OpenAI, UiPath, Anthropic, IBM, Automation Anywhere, ServiceNow	
	Foundation model labs: OpenAI, Anthropic, Google, Meta, xAI, Cohere, Mistral, Thinking Machines, Perplexity, Cognition, Cohere	
	Orchestration: AWS, Microsoft, Google, IBM, VMware, Oracle, Snowflake, Databricks, Datadog	Datacenter
	Datacenter: AWS, Microsoft, Google, Alibaba, Oracle, IBM, Tencent, Huawei, Meta, Apple, CoreWeave, Lambda	
	IT hardware: Dell, Cisco, HPE, Supermicro, Arista, NetApp, Foxconn, Quanta, Inventec, Inspur, Wiwynn, Lenovo, WD, Seagate	
	Power & thermal equipment: Schneider, Vertiv, Eaton, ABB, STULZ, Rittal, Caterpillar, Cummins, LiquidStack, Delta, nVent, Legrand	
	Real estate & colocation: Aligned, STACK, Vantage, QTS, CyrusOne, Applied Digital, Digital Realty, Equinix, CoreWeave, Lambda	Semiconductor
	Quantum computing: IBM, Google, Quantinuum, Microsoft, IonQ, PsiQuantum, AWS, Atom, Rigetti, Xanadu	
	Semiconductor chip design: NVIDIA, AMD, Intel, Google, AWS, Microsoft, Broadcom, Marvell, Meta, Cerebras, Groq, Tenstorrent, Hailo, Kneron, SambaNova, Axelera	
	Semiconductor supply chain: TSMC, ASML, SK hynix, Applied Materials, Lam Research, Synopsys, Cadence, ARM, Samsung, TEL, KLA	
Services: Accenture, IBM, Deloitte, PwC, EY, BCG, TCS, Capgemini, Wipro, Infosys		

Source: PitchBook



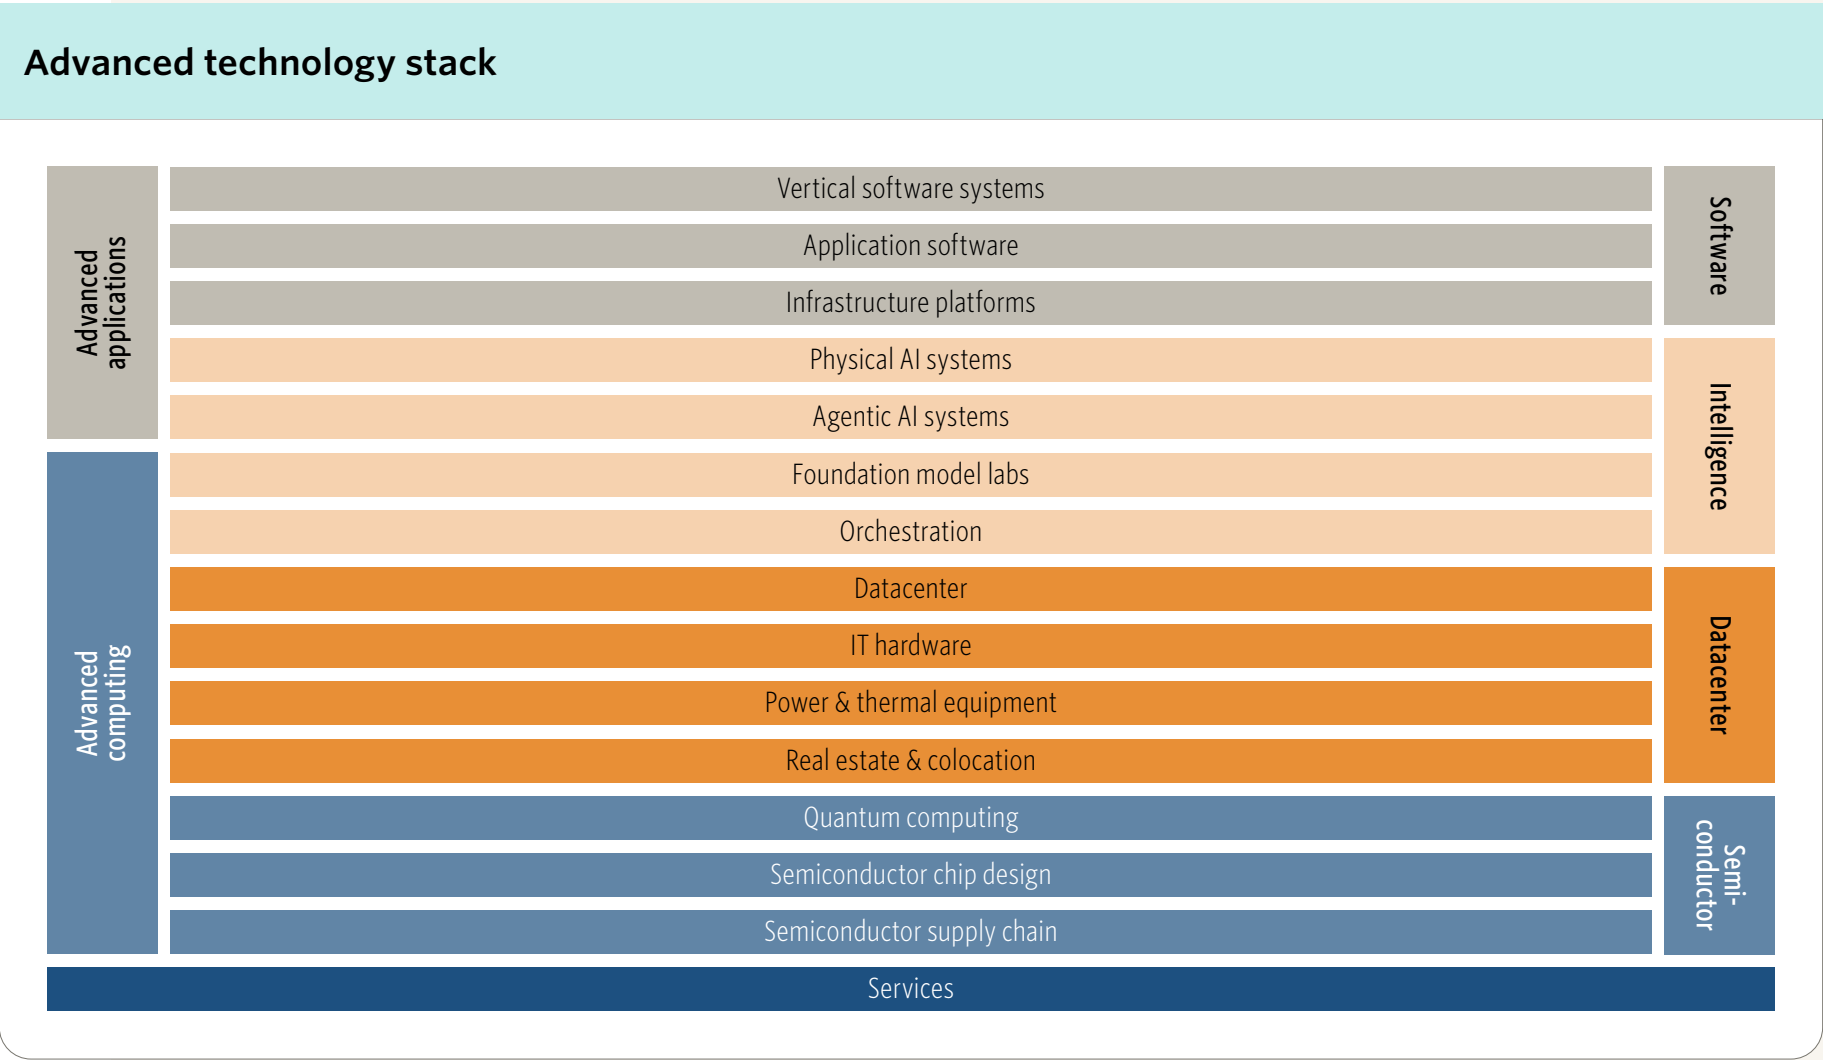
Overview and scope

Overview

Broadly defined, advanced computing includes all property, plants, equipment, and infrastructure that enables cloud computing at scale.

We divide advanced computing into seven segments across two markets: semiconductor computing and datacenter infrastructure. Within semiconductors, we track chip design, supply chains, and quantum computing. Within datacenters, we track facilities and power infrastructure, IT hardware and components, operators and service providers, and software and orchestration. In the [“Segment overview”](#) section of this report, the datacenter segments are examined as one group.

These seven segments are further divided into 40 categories. Segmentation is evenly split between semiconductors and datacenters. Advanced computing is a highly capital-intensive sector with extreme concentration in revenue, growth, financings, exits, and capital investments.



Source: PitchBook



OVERVIEW AND SCOPE

Given this market dynamic, we cover advanced computing across PE, VC, and public equity markets. We track deals, valuations, exits, and operating trends. The advanced computing [market map](#) will be updated periodically as categories are refined.

Please note that our advanced computing vertical does not track application layer segments. Vertical software systems, application software, infrastructure software, agentic AI, and physical AI are not included in hyperscaler capital expenditures. These applications are largely part of end-customer operating expenses, as is the case with software as a service (SaaS), or as production costs as embedded intelligence systems.

Language model companies, however, are critical capital investment catalysts for the AI datacenter ecosystem. Much of the increase in hyperscaler investments seen in 2025 was due to improvements in model fidelity and performance costs. We track language model labs and orchestration software categories in our vertical, as they are the prime software enablers of AI at scale.

Taxonomy scope

The advanced computing taxonomy is designed to be a comprehensive segmentation of all operational and information technologies used in the construction of AI datacenters. Operating expenses, such as SaaS, consulting, or utility costs, are not covered.

Within each segment, we follow all participating companies regardless of end-market exposure. For report coverage, however, we track only their advanced computing markets. For example, while Intel, Apple, and Qualcomm are de facto market leaders in AI client devices (PCs, mobile, wearables, and tablets), we track only their datacenter server or cloud computing end markets.

Market segmentation

- **Semiconductor supply chain:** This segment includes foundries, WFE, EDA and engineering software, IP, OSAT, materials, and advanced packaging.

- **Semiconductor chip design:** This segment includes graphic processors (GPUs), microprocessors (CPUs), AI accelerators, network processors, memory, analog mixed-signal chips, edge AI, power, optoelectronics, sensors, discretes, and passives.
- **Quantum computing:** This segment includes companies building hardware (processors and qubits), software (operating systems and algorithms), or cloud services specifically for quantum computing. Quantum error correction (QEC) and simulation are also included.
- **Datacenter facilities & power infrastructure:** This segment represents the physical real estate and utility layer of the datacenter. It includes buildings, power equipment, cooling and thermal management equipment, utility providers, and colocation management companies.
- **Datacenter IT hardware & components:** This segment includes computer, storage, and networking equipment. It ranges from fully integrated servers to components to cabling. Original equipment manufacturers (OEMs) and original design manufacturers (ODMs) are included.



OVERVIEW AND SCOPE

- **Datacenter operators & service providers:** This segment includes businesses operating datacenters to provide cloud services. Business models range from full stacks of on-demand computing (hyperscalers) to purpose-built AI factories to managed hosting services.
- **Datacenter software & orchestration:** This segment includes foundation models and software that manages, processes, automates, and extracts value from data infrastructure. This includes databases, hypervisors, data management, security, and monitoring.

In this segment, we track company partnerships, technology innovation, product introduction, and performance but not VC or PE activity. We view companies in this segment as enablers and catalysts of advanced computing capital expenditures but not directly tied to the build-out of AI datacenters.

About the data

Summary

The advanced computing report series tracks VC and PE deals globally. PitchBook clients have access to the complete analyst-

curated underlying data, which includes more than 2,000 companies segmented into 40 categories.

VC stage

The pre-seed/seed stage accounts for transactions held prior to Series A. “Early-stage VC” refers to Series A and B. “Late-stage VC” refers to Series C, D, and E. “Venture growth” refers to Series E+.

PE estimated deal count

The estimated deal count adjusts for normal lags in data collection. We arrive at our estimate by reviewing historical datasets. Due to lower deal counts, we do not provide this estimate at segment or category levels.

PE platform, add-on, and minority

- “Add-on” is any acquisition by a PE-backed company, regardless of target size.
- “Platform” is any buyout (majority equity acquisition) that is not an add-on, regardless of size.

- “Minority” refers to any minority equity investment, including follow-on investments by the same PE firms, and corresponds to the “PE growth/expansion” deal type in the PitchBook Platform.

Geographic scope

We chose a global scope for this report because the advanced computing vertical sees demand in virtually every geography, with the most developed markets in North America, Europe, South Korea, China, and Taiwan.

Investor segmentation

We track both VC and PE deal flow for the advanced computing vertical. While both semiconductor and datacenter operations are capital-intensive, the nature of their spending attracts different investors.

VCs dominate chip startup funding. Semiconductor startups are binary IP plays. Venture capital goes to fund R&D (at the early stages) and sales, marketing, and partnerships (at the later stages). These startups’ moat is design technology. The risk is that if the silicon fails or misses the market window,




OVERVIEW AND SCOPE

the venture equity is worthless. The return, however, is a significant stake in a successful company that has software-like gross margins (over 60%) and the potential for self-induced market dominance.

PE focuses on datacenters because they are real assets that produce yield. Essentially, datacenters are commercial real estate. Capital goes to land, concrete, power, cooling, and IT hardware. The moat is the physical infrastructure and customer contracts. The risk is execution, not technology. Datacenter operators structure take-or-pay or minimum-pay clauses to ensure they can recover capital investment costs. The return is a 10-to-15-year predictable revenue stream for PE.


To complete our coverage, we track the public company leaders in advanced computing. They dominate revenue and capital expenditure share and lead in design, manufacturing, software development, and foundation model intelligence.

Advanced computing VC deal count by quarter



Source: PitchBook • Geography: Global • As of September 30, 2025

Advanced computing PE deal count by quarter



Source: PitchBook • Geography: Global • As of September 30, 2025



Total addressable market

The advanced computing sector will be the beneficiary of the single largest deployment of investment capital in human history. AI, broadly defined, is the primary catalyst. Not since the invention of the assembly line for manufacturing, or widespread distribution of electricity into homes and businesses, have we seen the type of profound economic and societal change that AI can bring.

The TAM for this sector is hyperscaler capital expenditures. We forecast that 80% to 85% of all spending for advanced AI cloud computing will come from Microsoft, AWS, Google, Meta, and Oracle. These companies have established themselves as the AI computing platforms for enterprises worldwide, providing

them with nearly unlimited capital sources. We expect their dominance to remain for the foreseeable future.

The dual goal of AI is increased ROI for commercial businesses and increased quality of life for humanity. For cloud computing service providers, return on invested capital is the primary goal. For all other product companies, reduced production and operating expenses and increased customer stickiness are the goals.

The terms “superintelligence” and “artificial general intelligence” describe advanced reasoning capabilities of foundation models. As of this writing, foundation models had


Ph.D.-level knowledge in software coding, biology, physics, and chemistry. AI can now significantly enhance research and product development with no additional labor in these fields.

This is key to understanding the unprecedented increase in hyperscaler capital expenditures in 2025. Over the next five years, we anticipate at least \$6.4 trillion in investments in AI datacenters. Beginning this year, capital expenditure growth will exceed 45% per year for five years, with \$1.8 trillion spent in 2030 alone.



TOTAL ADDRESSABLE MARKET

Hyperscaler capital expenditures



Sources: Company filings and PitchBook • Geography: Global
As of November 26, 2025

Note: Data is based on PitchBook analyst calculations. The Big Five hyperscalers are Microsoft, Amazon, Google, Meta, and Oracle.



Growth drivers


Elasticity of intelligence

The goal of hyperscaler and foundation model investment is to provide expert-level knowledge to all levels of human enterprise for a monthly fee. This includes product manufacturing, R&D, sales and marketing, and general administrative tasks for commercial activities. For consumers, it includes access to a team of virtual experts in their pockets or houses working to make them productive, happy, and safe. The extent of these benefits is still unquantifiable.

We estimate that nearly \$4.7 trillion will be spent on IT infrastructure and \$1.2 trillion will be spent on power and cooling infrastructure over the next five years to accomplish the first phase of AI: expert-level intelligence.

The ARC-AGI-2 is a benchmark designed to test AI reasoning systems. As of this writing, Gemini 3 and GPT-5.1 (thinking, high) scored 32% and 18%, respectively. According to ARC Prize, the average human can score 60%, while a human expert can achieve 100%.¹ Any model achieving human-expert scores will fundamentally change the requirements of labor resources within every conceivable industry. This is the goal.

1: "ARC-AGI-2," ARC Prize, n.d., accessed November 26, 2025.



Sources: Company filings and PitchBook • Geography: Global • As of November 26, 2025
Note: Data is based on PitchBook analyst calculations.



GROWTH DRIVERS


Both Gemini and GPT achieved their scores at a cost of \$0.81 and \$1.17 per task, respectively.² This is enormously expensive when compared with GPT-5.1 (thinking, none) at \$0.058, which scored only 0.4%. In other words, to achieve superior intelligence, more high-cost computing is required. Equivalently, costs to mass-deploy superior intelligence will need to come down to the GPT 5.1 (thinking, none) cost levels after this high-performance computing infrastructure is significantly depreciated over time. However, both Gemini and GPT showed vast improvements in efficacy over prior versions due to architectural changes.

Crystallized intelligence

In January 2025, DeepSeek R1 was released with an MoE architecture and proved at scale that higher ARC-AGI-1 scores (basic information intelligence) were achievable at much lower costs. At the time, GPT-4.5 scored 10.3% at a cost of \$2.10, while R1 scored 15.3% at \$0.08. This was the “big bang” moment.

To achieve similar results, models were rearchitected into expert sections. No longer was an entire model needed to answer a prompt. Math, coding, creative writing, finance,

2: “ARC-AGI-2,” ARC Prize, n.d., accessed November 26, 2025.



Sources: Company filings, IDC, and PitchBook • Geography: Global • As of November 26, 2025
Note: Data is based on PitchBook analyst calculations.



GROWTH DRIVERS

science, and so on were separated from each other. The overall size of the model no longer affected the cost of operating the whole model; rather, only access to each expert system contributed to operating costs.

This allowed model builders to retrain expert sections with unconstrained parameter requirements because they knew that inference costs were going to be significantly lower. Foundation models quickly achieved expert-level skills in knowledge-based fields. This type of expertise is known as crystallized intelligence.

Models now routinely score highly in workflows such as software development, diagnostics, science, law, fraud detection, and customer service chats, offering enormous benefits to users. GPT-5.1 and Gemini 3 now rank 73% and 82% in ARC-AGI-1, respectively, and similarly in domain-specific science fields.

End users are now achieving vast increases in productivity from increased model expertise. Hyperscalers can now offer superior crystallized intelligence services at unit economics not possible prior to MoE architecture. The unprecedented revisions in growth of GPUs and networking in the datacenter market are a direct result of MoE rearchitecture. We forecast \$2.8 trillion in capital investment in datacenter processors alone over the next five years. By 2030, nearly 70% of all processors will be in the datacenter market for training, inference, and networking.

Fluid intelligence

The lower cost per task has also allowed model operators to focus on reasoning. We define LMF as the ability of a language model to understand and execute what the model user intends. This requires not just crystallized intelligence but also reasoning. Reasoning, known as fluid intelligence, is the ability

to solve novel problems by drawing logical conclusions from limited data. With lower costs per task, foundation models now use additional computing resources to reason through more complex user prompts.

When a model thinks, internal tokens are generated to break down problems, evaluate different approaches, explore different solutions, check work, and backtrack if needed. The result is a more accurate and useful inference. The ARC-AGI-2 scores of GPT and Gemini are vastly improving but at a very high cost. This is the elasticity of intelligence: More computing generates more intelligence.

To achieve expert-level fluid intelligence, models will be architected to think longer, requiring more computing. The investment implication is staggering: Intelligence scales with computing power. Demand for GPUs is not capped by the number of users but by the complexity of the problems that




GROWTH DRIVERS

users solve. Searching on Google servers requires a certain level of investment to maintain high efficacy and strong profitability. Advanced solutions offered by AI experts (such as generative applications, agentic AI, and physical AI) will require an entirely different level of capital investment.

Power and cooling

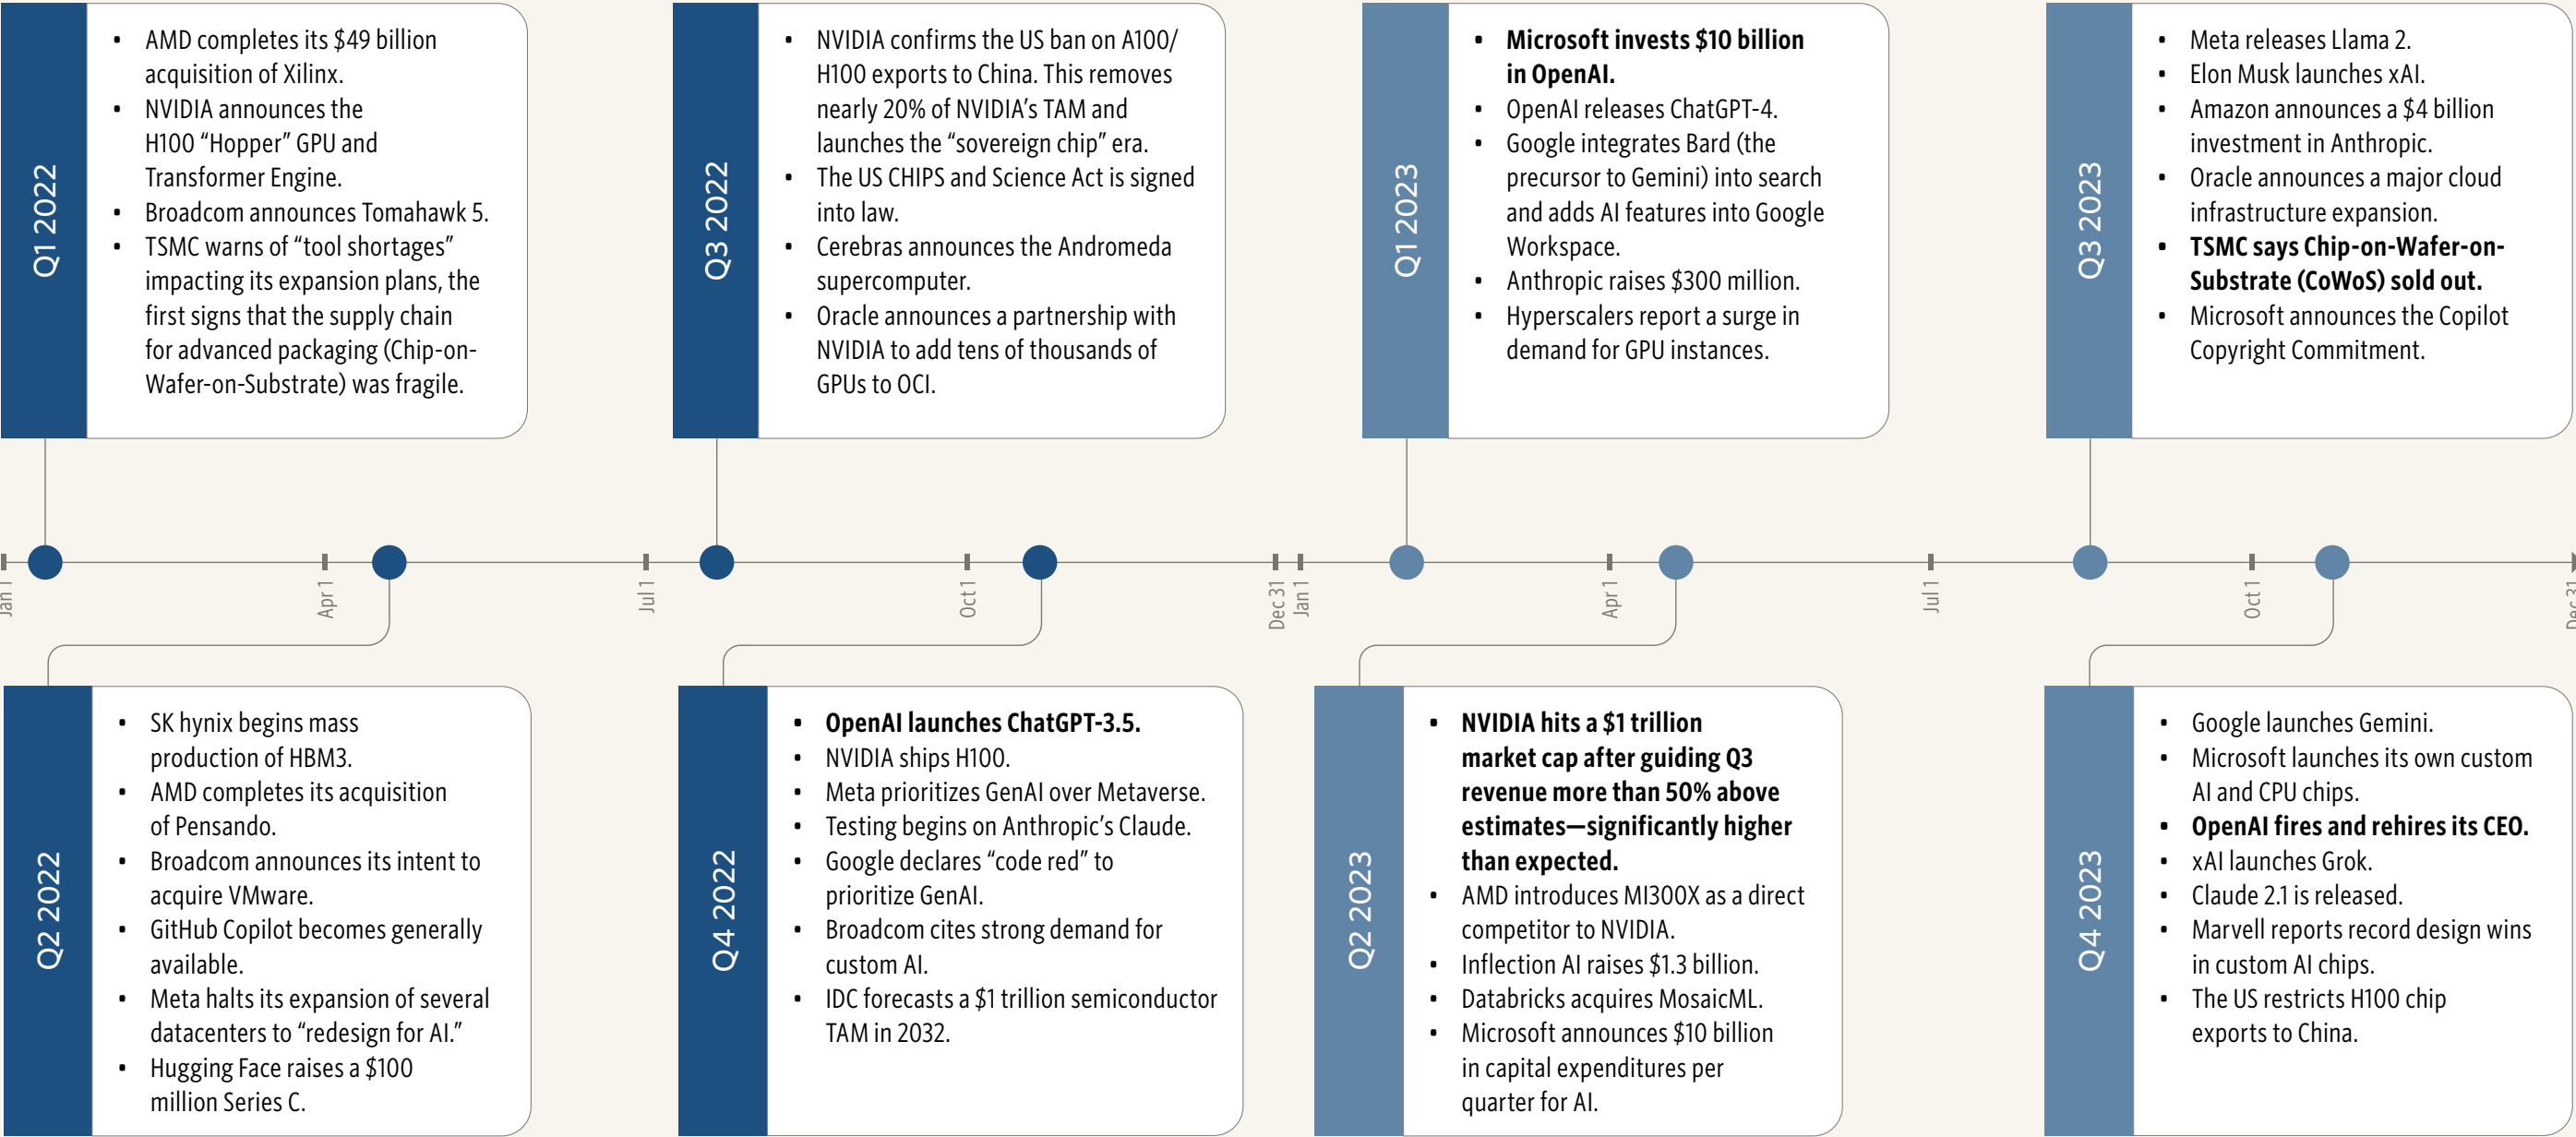
The seemingly unreasonable power and computing requirement needed beyond 2030, as championed by OpenAI, is not fantasy. Power and cooling are the gating factors to investment and require management execution that very few companies have. We see a CAGR of nearly 34% for power and cooling equipment through 2030, faster than all other segments. Power and cooling are the most important enablers of growth, provided these companies can produce equipment at rates faster than the growth of hyperscaler capital expenditures.



Sources: Company filings and PitchBook • Geography: Global • As of November 26, 2025
Note: Data is based on PitchBook analyst calculations.

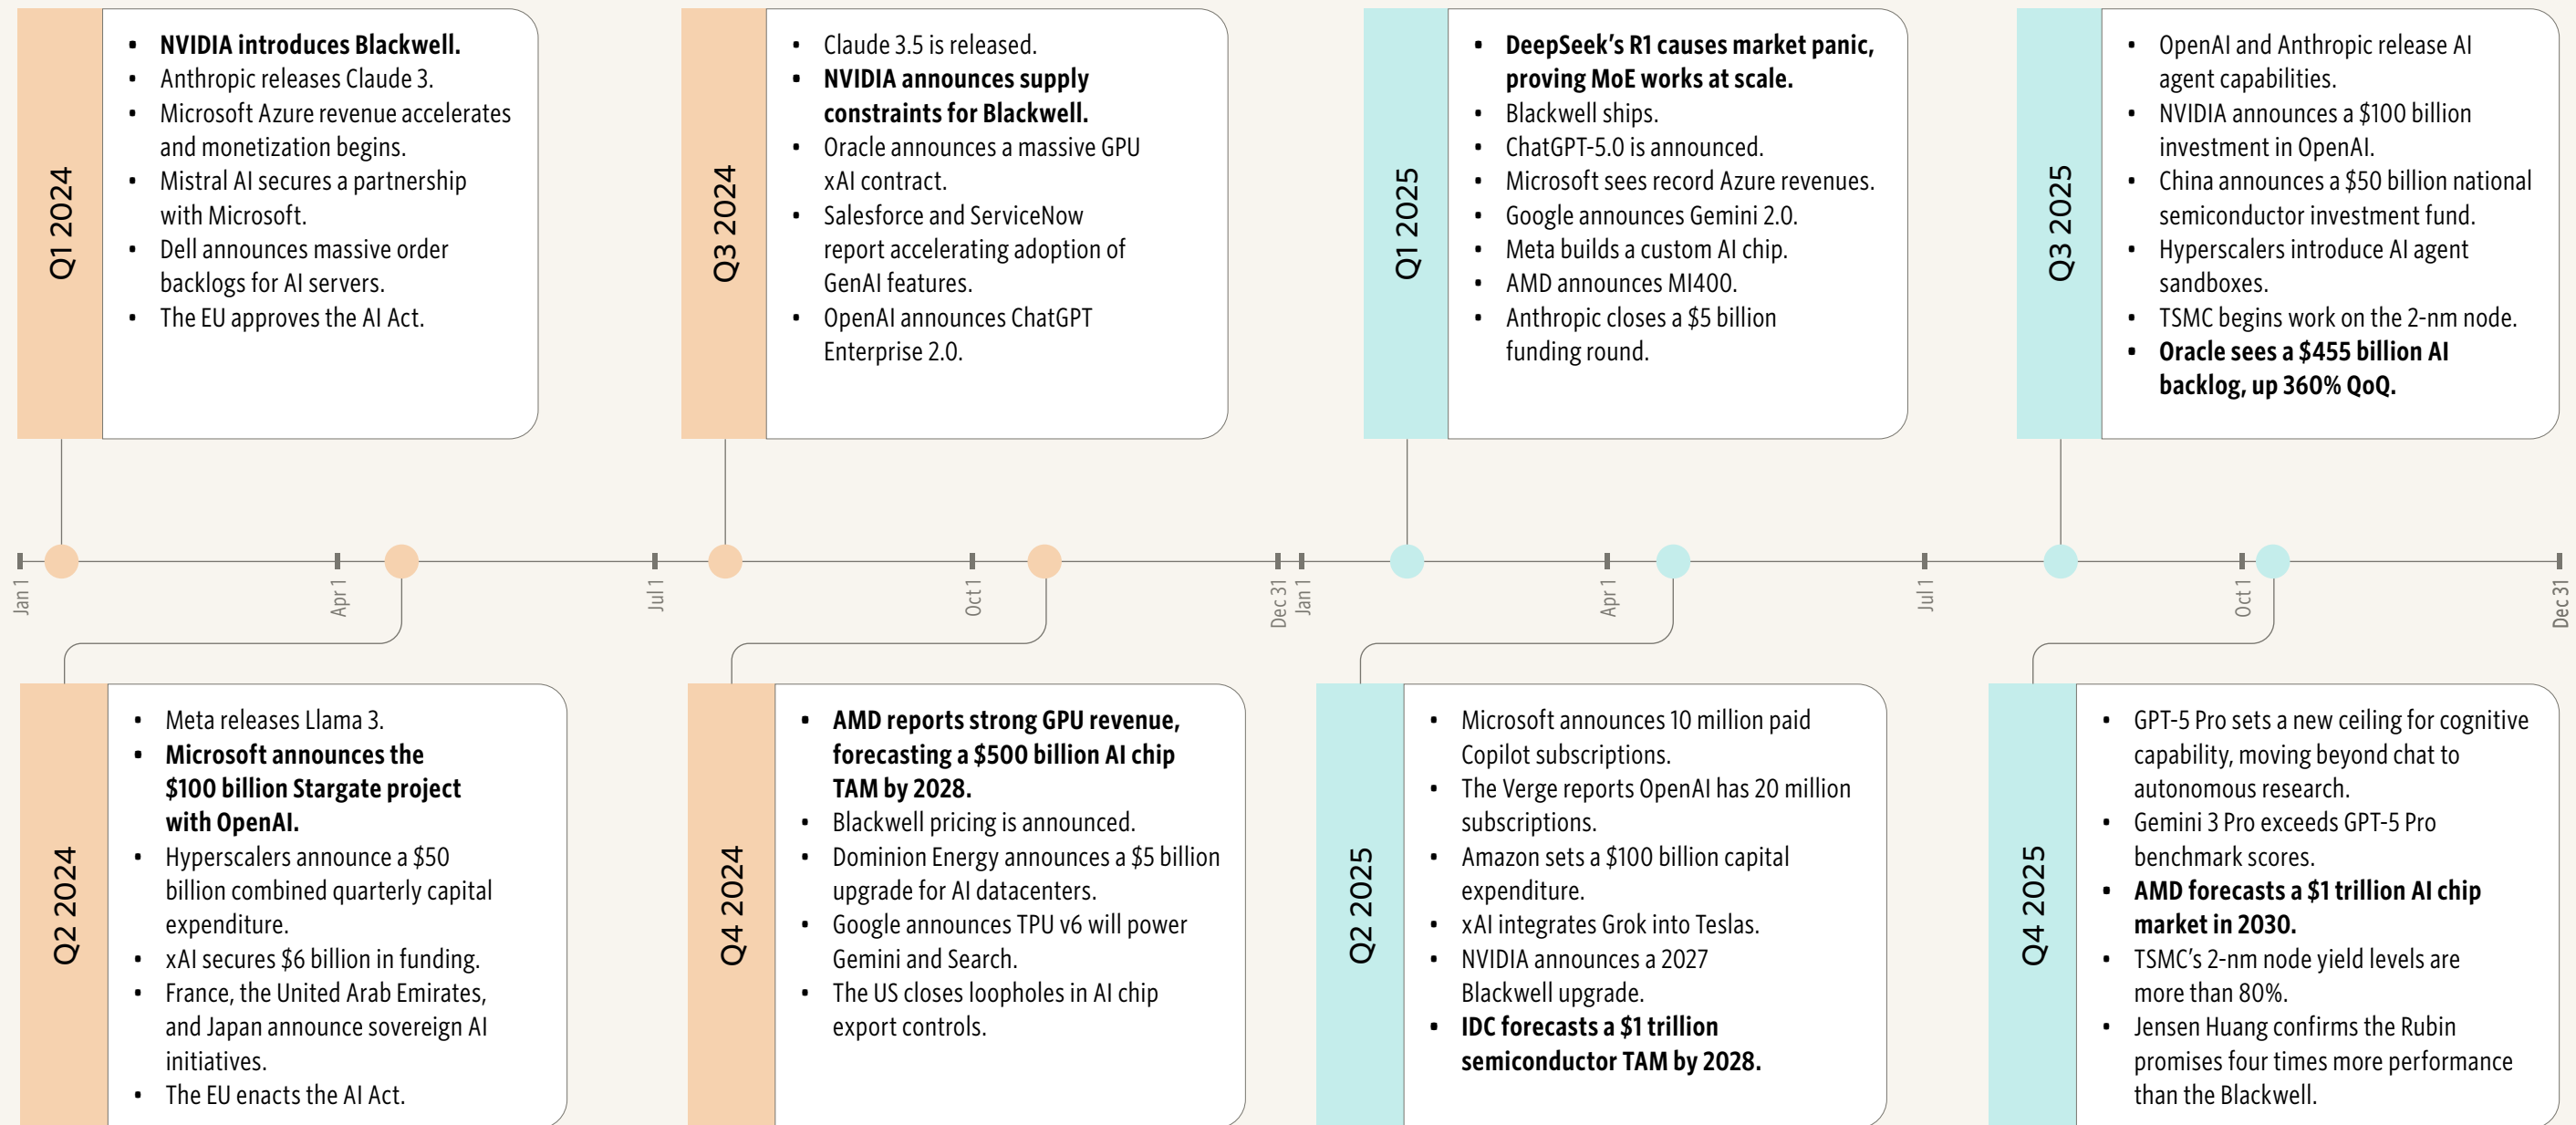


Pivots and catalysts






PIVOTS AND CATALYSTS





PIVOTS AND CATALYSTS

NVIDIA market capitalization and quarterly revenue (\$B)



Source: Company filings • Geography: Global • As of November 26, 2025



Leadership matrix

We developed the following leadership matrix to quantify company dominance within advanced computing categories. We prioritize AI criticality over market share. AI criticality considers how essential a company is to the advancement of the AI ecosystem and its irreplaceability within the technology stack. At the time of this writing, we had evaluated 539 public and private companies. This count will expand with each new report.

In determining the leadership score, we evaluate market dominance, market importance, technology leadership, and moat defensibility. A high score (100 to 80) indicates the company owns a strategic platform or dominates essential IP and is dominant in its end markets. A medium score (50 to 79) indicates a significant player but one that does not own a technology monopoly. A low score (2 to 49) indicates an ecosystem participant that is either a niche player, a commoditized supplier, or a secondary competitor.

We score companies against our taxonomy of seven segments and 40 categories. This prevents scores based solely on primary business lines and ensures accurate peer comparisons within each category.

Semiconductor computing

Semiconductor supply chain

Scores in this category derive from process-technology criticality and manufacturing capacity. Monopoly-tier enablers such as lithography providers receive the highest ratings. Commoditized packaging or assembly players receive lower scores. We evaluate the difficulty of replacing the vendor in times of shortage.

Semiconductor chip design

We evaluate market-share dominance in specific AI workloads. Companies designing custom silicon for training and inference receive the highest scores. General-purpose chipmakers score lower than AI-specific designers. We focus on the chips that run AI datacenters.

Quantum computing

We assess technological leadership in qubit maturity and error correction. This distinguishes viable hardware from theoretical research. We value operational systems over future promises. This identifies leaders in the next generation of computing.

Datacenter infrastructure

Facilities & power infrastructure

- **Real estate & colocation:** We evaluate operational square footage and megawatt capacity under management. This bucket includes both landlords owning shells and builders constructing them. It strictly measures physical infrastructure capacity. We exclude cloud tenants to avoid double counting.
- **Power utilities:** We measure regulated utilities on gigawatts of owned generation capacity. Control over the actual electrons is the primary metric. We exclude entities that merely procure power from the grid. This isolates the producers from the consumers.
- **Power equipment:** Scores reward manufacturers of critical hardware, such as generators and fuel cells. We focus on the OEMs building the physical engines. Facilities that simply install these systems are excluded. This separates technology creators from technology users.
- **Cooling & thermal equipment:** Leadership is defined by high-density thermal management technologies. We



LEADERSHIP MATRIX

prioritize liquid cooling, immersion, and direct-to-chip solutions. Standard HVAC providers score lower than AI-specific thermal experts. This reflects the thermal demands of GPUs and AI accelerators.

IT hardware & components

We evaluate global shipment volumes of AI-optimized systems. We score integrators building physical servers, racks, storage, and switches highest. This captures the physical IT infrastructure of the datacenter.

Operators & service providers

This includes all hyperscalers and cloud computing service providers. We look beyond general cloud revenue to specific AI computing capacity. This approach rewards AI density (processing power per square foot) and massive GPU

availability. Specialized neoclouds score comparably to traditional hyperscalers in computing performance per watt but lack computing capacity.




Software & orchestration

- **Foundation models:** Scores reflect performance benchmarks and adoption of proprietary models. We focus on labs owning the weights, not just providing an API. Ownership of core IP is the deciding factor. This isolates true model innovation from simple integration.
- **Orchestration:** We evaluate software developer adoption of these tools. This includes vector databases, data lakes, and machine learning operations (MLOps) platforms. We gauge the stickiness of the development stack. This identifies the essential tools for building AI products.












LEADERSHIP MATRIX: SEMICONDUCTOR COMPUTING

Semiconductor supply chain

Company		Rank
TSMC		100
ASML		99
KLA Corporation		96
Tokyo Electron		96
Lam Research		95
Applied Materials		95
SiFive		88

Source: PitchBook • Geography: Global • As of November 30, 2025
Note: The top rankings are cut off at 88.

Semiconductor chip design





Company		Rank
NVIDIA		100
Google (Alphabet)		99
SK hynix		95
Apple		95
Samsung		94
Cerebras Systems		93
Microsoft		93
AMD		92
Micron Technology		91

Source: PitchBook • Geography: Global • As of November 30, 2025
Note: The top rankings are cut off at 91.



LEADERSHIP MATRIX: SEMICONDUCTOR COMPUTING

Quantum computing

















Company		Rank
Classiq		100
Riverlane		100
Infleqtion		100
IQM Quantum Computers		100
Pasqal		100
PsiQuantum		100
QuEra Computing		100
Xanadu		100
Alice & Bob		100

Source: PitchBook • Geography: Global • As of November 30, 2025
Note: The top rankings are cut off at 100.



















LEADERSHIP MATRIX: DATACENTER INFRASTRUCTURE

Facilities & power infrastructure

Real estate & colocation					
Real estate company		Rank	Power utility company		Rank
Aligned Data Centers		100	Constellation Energy		100
Brookfield		100	NextEra Energy		100
CloudHQ		100	American Electric Power		98
CyrusOne		100	BGE		98
DataBank		100	Con Edison		98
Digital Realty		100	Dominion Energy		98
Equinix		100	Duke Energy		98
Flexential		100	Exelon		98

Source: PitchBook • Geography: Global • As of November 30, 2025.
Note: The top rankings for real estate and power utility companies are cut off at 100 and 98, respectively. Not all ties are shown.

Facilities & power infrastructure (continued)

Power & thermal equipment					
Power equipment company		Rank	Cooling/thermal equipment company		Rank
ABB Ltd		100	Vertiv		100
Vertiv		98	Eaton		91
Eaton		98	Schneider Electric		91
Bloom Energy		98	Asetek		91
Caterpillar		98	Bosch		91
Cummins		98	Boyd		91
Enphase		98	Carrier		91
Eos Energy		98	Chatsworth Products		91

Source: PitchBook • Geography: Global • As of November 30, 2025
Note: The top rankings for power equipment and cooling/thermal equipment companies are cut off at 98 and 91, respectively.
Not all ties are shown.









LEADERSHIP MATRIX: DATACENTER INFRASTRUCTURE

IT hardware & components

IT hardware		
Company		Rank
Supermicro		100
Wiwynn		100
Arista Networks		98
ASUSTeK Computer		98
Foxconn		98
Inventec		98
Lenovo		98
Quanta Cloud Technology (QCT)		98

Source: PitchBook • Geography: Global • As of November 30, 2025
Note: The top rankings are cut off at 98.

Operators & service providers







Datacenter services		
Company		Rank
Amazon (AWS)	AWS China physical infrastructure assets	100
Google (Alphabet)		100
Microsoft		100
Alibaba		95
Huawei		90
Baidu		90
Tencent		90

Source: PitchBook • Geography: Global • As of November 30, 2025
Note: The top rankings are cut off at 90.








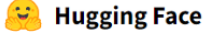

LEADERSHIP MATRIX: DATACENTER INFRASTRUCTURE

Software & orchestration

Foundation models		
Company		Rank
OpenAI		100
Google (Alphabet)		100
Anthropic		98
Microsoft		95
Meta		95
Amazon (AWS)	AWS China physical infrastructure assets	92
Mistral AI		92

Source: PitchBook • Geography: Global • As of November 30, 2025
Note: The top rankings are cut off at 92.

Software & orchestration (continued)

Orchestration		
Company		Rank
Microsoft		100
Google (Alphabet)		98
Amazon (AWS)	AWS China physical infrastructure assets	98
Apple		96
Databricks		96
NVIDIA		96
Hugging Face		91
Snowflake		91

Source: PitchBook • Geography: Global • As of November 30, 2025
Note: The top rankings are cut off at 91.



Segment overview

Semiconductor supply chain


This segment includes foundries, WFE, EDA and engineering software, IP, OSAT, materials, and advanced packaging.

The supply chain segment’s revenues are reflected in the cost of goods sold for both integrated device manufacturers (IDMs) and fabless semiconductor companies. The supply chain is a single global network, with Asia-Pacific producing most of the world’s memory and leading-edge processors and accounting for the majority of chip testing and assembly. Supply chain sovereignty will take at least a decade to adjust geographic concentration.

We see foundries and WFE as the prime market-share beneficiaries of hyperscaler semiconductor purchasing. Leading-edge process nodes are needed to produce high-performance GPU, CPU, high-bandwidth memory (HBM), and AI accelerators.

We see advanced packaging, IP, and EDA as leading growth within the supply chain segment. Advanced processors are moving toward lower-power ARM and RISC-V designs.

Semiconductor supply chain TAM



Sources: Company filings and PitchBook • Geography: Global • As of November 26, 2025
Note: Data is based on PitchBook analyst calculations.



SEGMENT OVERVIEW

We forecast \$3.1 trillion in cumulative semiconductor supply chain revenues with a five-year CAGR of 8.6%. EDA/IP and foundries are forecast to grow 12.2% and 10.5%, respectively.

Semiconductor chip design

This segment includes graphic processors (GPUs), microprocessors (CPUs), AI accelerators, network processors, memory, analog mixed-signal chips, edge AI, power, optoelectronics, sensors, discretes, and passives.

The semiconductor chip design segment comprises IDMs and fabless semiconductor companies. This has traditionally been a highly cyclical business with booms and crashes every three years. We forecast \$3.7 trillion in industry revenues over the next five years with no cyclicalities in the forecast period. This is due to the unprecedented adoption of AI processors. We forecast the total datacenter semiconductor five-year CAGR to be 25.6%. This segment will increase from 41% to 63% of the total semiconductor market.

By 2030, we expect GPUs, CPUs, and AI accelerators to represent 57.1% of total datacenter silicon. We see networking processors at 17.1%, HBM at 14.5%, and rack power ICs at 4.8%.

As shown in the supply chain TAM chart, we expect gross margins to improve significantly over the next five years. The cost to produce high-performance semiconductors for AI applications is significantly less than that of traditional processors. GPU manufacturers utilize hardware-software co-design for AI and high-performance computing platforms.

This approach involves the simultaneous, iterative design of specialized hardware (GPUs, accelerators, and networking) and optimized software (CUDA, ROCm, libraries, and frameworks) to achieve maximum performance/cost. This added integration commands higher gross margins, making chip manufacturing costs lower as a percentage of the final selling prices. Higher gross margins for AI processors and AI memory can provide significant upside to our \$1.6 trillion TAM forecast for 2030.

Quantum computing

We believe widespread commercial viability for quantum computing is unlikely before 2035. This is due to the immense engineering hurdles of QEC and the concurrent rapid advancement of AI on accelerated classical hardware. While current noisy intermediate-scale quantum devices exist, they lack the fault tolerance required for commercial applications.

Qubits are fragile. Maintaining their quantum state (coherence) requires isolation from environmental noise. To make them useful, QEC is needed to turn thousands of physical qubits into just one stable logical qubit. This creates massive overhead. Building a quantum computer with millions of physical qubits to get a few thousand logical qubits will likely require a decade of additional engineering development.

The rise of AI and high-performance computing has reduced quantum computing's TAM. Many problems originally targeted by quantum computing, such as complex optimization, protein folding, and material science simulations, are currently being addressed by advanced AI models running on classical supercomputers (GPUs). For the next decade, this "good enough" computing, driven by AI's massive capital injection, will likely be sufficient for most industrial tasks.

IBM, Google, and Microsoft are leading quantum research. IBM is aggressively ramping its road map toward "quantum utility," focusing on scaling its superconducting processors to thousands of qubits while integrating error-mitigation techniques. Google Quantum AI is prioritizing low-error logical qubits, recently proving that increasing the number of physical qubits can suppress errors in a logical qubit.




SEGMENT OVERVIEW

Microsoft is making a longer-term bet on “topological qubits,” a theoretically more stable but harder-to-engineer design, aiming for a fully fault-tolerant machine from the start.

Companies such as IonQ and Quantinuum are commercializing trapped-ion systems that have higher fidelity and connectivity than superconductors, offering cloud access to researchers. Software-focused startups are developing quantum-inspired algorithms that run on classical hardware to deliver immediate value to enterprise clients while preparing them for future quantum hardware.


By 2035, we expect the industry to transition from experimental noise-prone systems to the first generation of fault-tolerant quantum computers. Successful error correction will unlock specific high-value commercial applications. We expect IBM, Google, and Microsoft to be the sources of innovation for the foreseeable future.

Semiconductor chip design TAM



Sources: Company filings, IDC, and PitchBook • Geography: Global • As of November 26, 2025
Note: Data is based on PitchBook analyst calculations.

Quantum computing VC deal activity by quarter




Source: PitchBook • Geography: Global • As of September 30, 2025




SEGMENT OVERVIEW

Quarterly semiconductor computing VC deal activity by segment




Source: PitchBook • Geography: Global • As of September 30, 2025

Quarterly semiconductor computing PE deal activity by segment




Source: PitchBook • Geography: Global • As of September 30, 2025

Quarterly semiconductor computing VC exit activity by segment



Source: PitchBook • Geography: Global • As of September 30, 2025

Quarterly semiconductor computing PE exit activity by segment



Source: PitchBook • Geography: Global • As of September 30, 2025



SEGMENT OVERVIEW

Datacenter infrastructure

The datacenter infrastructure market includes all property, plants, and equipment needed to build out AI computing to scale. We divide this market into four segments: facilities & power infrastructure, IT hardware & components, operators & service providers, and software & orchestration.

The AI build-out represents a generational capital deployment cycle of unprecedented scale. Given the magnitude of capital expenditure, we see access to power as the primary growth enabler and driver within advanced computing.

Within the facilities & power infrastructure segment, we track power-generation innovation and new construction of facilities. Site selection is now focused on secure power access over proximity to population centers or fiber-optic connections. Datacenter operators seek regions with low cost, renewable energy, and favorable grid interconnection queues. Time to power is the gating factor as building design specs are shifting quickly to accommodate extreme power densities.

For reference, traditional datacenters support 10 to 15 kW per rack. Advanced AI clusters require 50 to 100 kW per rack,

requiring reinforced floors and specialized liquid-cooling loops. AI workloads rely on high-performance GPUs and accelerators that consume significantly more power than traditional CPU servers. This is pushing power densities far beyond legacy capabilities, requiring new solutions.

Power equipment OEMs are critical for on-site generation, power availability, and redundancy. Natural gas and diesel generators remain standard for primary and backup power. Cooling and thermal management equipment is no longer a secondary function but a primary requirement. High-density GPUs require direct-to-chip or immersion cooling to prevent thermal throttling. This shift disrupts legacy air-cooling vendors for liquid-cooling specialists.

Within the IT hardware & components segment, we track datacenter ICs: processors, accelerators, HBM, and networking ICs. Semiconductor companies do not sell directly to hyperscalers but rather work with ODMs. Datacenters require customized computing specs that are far more powerful than specs for traditional enterprises. ODMs build systems that require specialized chassis designs to manage heavy components and extreme heat. Storage systems have also advanced to store data and serve it to GPUs without latency

using all-flash arrays. Data throughput speed is now as critical as total storage capacity for training workloads.

We see the networking hardware market as the fastest growth market within datacenter infrastructure. Optical interconnects and transceivers are vital for linking thousands of GPUs into a single cluster. Bandwidth limitations stall millions of dollars of revenue, necessitating a massive near-term upgrade cycle.

Within the operators & service providers segment, we track the AI cloud service providers (CSPs) themselves. CSPs can be divided into generalist hyperscalers and specialized neoclouds offering bare-metal GPU access. Neoclouds are a new market. Given the vast demand of AI computing and secure power, neoclouds are able to fill the time-to-power gap in the near term.

Hyperscalers leverage massive scale while neoclouds compete on availability of GPUs. We forecast that the Big Five hyperscalers (Microsoft, Amazon, Google, Meta, and Oracle) will account for 80% to 85% of worldwide capital expenditures on advanced datacenters for the foreseeable future. Their spending is limited by operating cash flow and customer productivity gains.



SEGMENT OVERVIEW


Foundation model companies rely heavily on hyperscaler computing platforms to offer expert-level AI to their professional and enterprise customers. Hyperscalers, similarly, partner with foundation models to run these workloads at ever-higher throughputs and uptime rates.

As discussed previously, AI is significantly elastic at this point in its development. The more computing power AI is given, the more intelligent and useful it becomes. The goal for a foundation model is perfect fidelity with expert intelligence. When this is achieved, agentic and physical AI will replace entry- and intermediate-level human skills, requiring even more computing. Usefulness will be defined by labor substitution within the next five years.

As a result, we track operational and performance trends at every service level. Currently, model usefulness is outpacing capital expenditure deployment. This is creating a demand imbalance and is accelerating revenue growth to levels never seen before.

Within the software & orchestration segment, we track the bleeding edge of AI innovation. This includes foundation model companies and AI datacenter infrastructure tools. Unlike all other segments within advanced computing, companies in

Hyperscaler free cash flow



Sources: Company filings and PitchBook • Geography: Global • As of November 26, 2025
Note: Data is based on PitchBook analyst calculations. Hyperscaler operating cash flow is derived from all company revenues, not just cloud services.



SEGMENT OVERVIEW

this segment are subject to the operating expense budget of the cloud operator or the end customer. Efficiency at this layer directly translates to lower total cost of ownership.

Infrastructure platforms such as data warehouses and data lakes are critical for managing the massive datasets required for training, analytics, and business intelligence. They enable RAG, connecting models to secure enterprise data. Data security and governance are also tracked in this segment.

We also track MLOps platforms that manage AI model lifecycles. They ensure reproducible training runs and efficient resource utilization. This software infrastructure is becoming as essential as the hardware itself. MLOps are crucial for

enterprises, as they provide a framework for developing, deploying, and managing AI solutions across various departments and business units.

Legacy virtualization and management software is also tracked within this segment. We see virtualization as way for foundation model companies to gain a foothold within the enterprise. Consumer-facing chatbots are a commodity. We see the rise of expert intelligence as a key growth driver to enterprise AI adoption. We predict foundation model companies will begin to partner with and ultimate acquire enterprise infrastructure SaaS companies to maintain growth of their customer bases and enhance account control.


We see upside to our \$1.7 trillion datacenter capital expenditure estimate for 2030. While we view power as the gating factor, we believe that expert intelligence is growing faster than anticipated, increasing the need for more capital expenditure. The introduction of Gemini 3 Pro at an ARC-AGI-2 score of 32% is an extraordinary accomplishment. The race to model expert intelligence is here. This will dramatically increase the usefulness and profitability of datacenter revenue.

We see 2029 as the breakout year for hyperscaler free cash flow as core business margins accelerate due to expert intelligence systems simultaneously optimizing internal costs and product revenue demand catalysts. This will be the dry powder needed to invest in the next wave of advanced computing: agentic and physical AI.




SEGMENT OVERVIEW

Quarterly datacenter infrastructure VC deal activity by segment




Source: PitchBook • Geography: Global • As of September 30, 2025

Quarterly datacenter infrastructure PE deal activity by segment




Source: PitchBook • Geography: Global • As of September 30, 2025

Quarterly datacenter infrastructure VC exit activity by segment



Source: PitchBook • Geography: Global • As of September 30, 2025

Quarterly datacenter infrastructure PE exit activity by segment




Source: PitchBook • Geography: Global • As of September 30, 2025




Deal and exit activity

Advanced computing VC deal value (\$B) by quarter




Source: PitchBook • Geography: Global • As of September 30, 2025

Advanced computing VC exit value (\$B) by quarter




Source: PitchBook • Geography: Global • As of September 30, 2025

Advanced computing PE deal value (\$B) by quarter



Source: PitchBook • Geography: Global • As of September 30, 2025

Advanced computing PE exit value (\$B) by quarter




Source: PitchBook • Geography: Global • As of September 30, 2025




DEAL AND EXIT ACTIVITY

Semiconductor computing VC deal and exit count by quarter



Source: PitchBook • Geography: Global • As of September 30, 2025

Semiconductor computing PE deal and exit count by quarter




Source: PitchBook • Geography: Global • As of September 30, 2025




DEAL AND EXIT ACTIVITY

Datacenter infrastructure VC deal and exit count by quarter



Source: PitchBook • Geography: Global • As of September 30, 2025

Datacenter infrastructure PE deal and exit count by quarter



Source: PitchBook • Geography: Global • As of September 30, 2025



DEAL AND EXIT ACTIVITY

Top 20 datacenter infrastructure companies by VC deal value (\$B)

Company	Segment	2020 deal value	2021 deal value	2022 deal value	2023 deal value	2024 deal value	2025 deal value	2020-2025 total deal value
DayOne Systems	Facilities & power infrastructure	N/A	N/A	N/A	\$0.7	\$1.2	N/A	\$1.9
CoreWeave	Operators & service providers	N/A	\$0.1	N/A	\$0.4	\$1.1	N/A	\$1.6
JD Property	Facilities & power infrastructure	N/A	\$0.7	\$0.8	N/A	N/A	N/A	\$1.5
Crusoe	Operators & service providers	N/A	\$0.1	\$0.5	N/A	\$0.7	N/A	\$1.3
GreenScale	Facilities & power infrastructure	N/A	N/A	N/A	N/A	\$1.3	N/A	\$1.3
Princeton Digital Group	Facilities & power infrastructure	\$0.4	N/A	\$0.5	N/A	N/A	N/A	\$0.9
Bohao Internet Data Services	Facilities & power infrastructure	\$0.3	N/A	\$0.5	N/A	N/A	N/A	\$0.7
KT Cloud	Operators & service providers	N/A	N/A	N/A	\$0.5	N/A	N/A	\$0.5
Compute North	Facilities & power infrastructure	N/A	\$0.0	\$0.4	N/A	N/A	N/A	\$0.4
Adept	Software & orchestration	N/A	N/A	\$0.1	\$0.4	N/A	N/A	\$0.4

Source: PitchBook • Geography: Global • As of September 30, 2025



DEAL AND EXIT ACTIVITY

Top 20 datacenter infrastructure companies by VC deal value (\$B) (continued)

Company	Segment	2020 deal value	2021 deal value	2022 deal value	2023 deal value	2024 deal value	2025 deal value	2020-2025 total deal value
Aiven	Software & orchestration	\$0.0	\$0.2	\$0.2	N/A	N/A	N/A	\$0.4
AVAIO Digital	Facilities & power infrastructure	N/A	\$0.4	N/A	N/A	N/A	N/A	\$0.4
Fermi (REIT)	Facilities & power infrastructure	N/A	N/A	N/A	N/A	N/A	\$0.4	\$0.4
Natron Energy	Facilities & power infrastructure	\$0.0	N/A	\$0.1	N/A	\$0.2	\$0.1	\$0.3
Vultr	Operators & service providers	N/A	N/A	N/A	N/A	\$0.3	N/A	\$0.3
Evolution Data Centres	Facilities & power infrastructure	N/A	N/A	N/A	N/A	N/A	\$0.3	\$0.3
Anyscale	Software & orchestration	\$0.1	N/A	\$0.2	N/A	N/A	N/A	\$0.3
Sakana AI	Software & orchestration	N/A	N/A	N/A	N/A	\$0.2	N/A	\$0.2
Kao Data	Facilities & power infrastructure	N/A	\$0.2	N/A	\$0.1	N/A	N/A	\$0.2
City Cloud International	Operators & service providers	\$0.1	\$0.2	N/A	N/A	N/A	N/A	\$0.2

Source: PitchBook • Geography: Global • As of September 30, 2025



DEAL AND EXIT ACTIVITY

Top 20 semiconductor computing companies by VC deal value (\$B)

Company	Segment	2020 deal value	2021 deal value	2022 deal value	2023 deal value	2024 deal value	2025 deal value	2020-2025 total deal value
ChangXin Memory Technologies	Semiconductor chip design	\$2.4	\$0.1	\$1.3	N/A	\$1.5	N/A	\$5.3
Groq	Semiconductor chip design	N/A	\$0.3	N/A	N/A	\$0.7	\$2.3	\$3.2
Changxin Xinqiao	Semiconductor chip design	N/A	N/A	N/A	\$2.0	\$1.2	N/A	\$3.2
GTA Semiconductor	Semiconductor supply chain	N/A	\$1.3	N/A	\$1.9	N/A	N/A	\$3.1
MetaX Tech	Semiconductor chip design	\$0.0	\$0.2	\$0.1	\$0.1	\$0.1	\$1.2	\$1.7
ESWIN Material	Semiconductor supply chain	N/A	\$0.5	\$0.8	\$0.3	N/A	N/A	\$1.7
PsiQuantum	Quantum computing	N/A	\$0.5	N/A	N/A	N/A	\$1.0	\$1.5
Horizon Robotics	Semiconductor chip design	\$0.2	\$1.1	N/A	\$0.2	\$0.0	N/A	\$1.4
Cerebras Systems	Semiconductor chip design	N/A	\$0.2	N/A	N/A	N/A	\$1.1	\$1.3
Ferrotec	Semiconductor supply chain	\$0.6	\$0.5	\$0.2	N/A	N/A	N/A	\$1.3

Source: PitchBook • Geography: Global • As of September 30, 2025



DEAL AND EXIT ACTIVITY

Top 20 semiconductor computing companies by VC deal value (\$B) (continued)

Company	Segment	2020 deal value	2021 deal value	2022 deal value	2023 deal value	2024 deal value	2025 deal value	2020-2025 total deal value
ESWIN Computing	Semiconductor chip design	\$0.4	\$0.4	N/A	\$0.4	N/A	\$0.0	\$1.2
Chipone Technology	Semiconductor chip design	\$0.1	\$1.0	N/A	N/A	N/A	N/A	\$1.1
Tenstorrent	Semiconductor chip design	N/A	\$0.2	N/A	\$0.1	\$0.7	N/A	\$1.0
Biren Technology	Semiconductor chip design	\$0.4	\$0.2	\$0.0	\$0.1	N/A	\$0.2	\$1.0
SambaNova Systems	Semiconductor chip design	\$0.3	\$0.7	N/A	N/A	N/A	N/A	\$0.9
Quantinuum	Quantum computing	N/A	N/A	\$0.0	N/A	\$0.3	\$0.6	\$0.9
Silicon Magic Semiconductor	Semiconductor chip design	\$0.5	N/A	\$0.4	N/A	N/A	N/A	\$0.9
Lightmatter	Semiconductor chip design	N/A	\$0.1	N/A	\$0.3	\$0.4	N/A	\$0.8
Hua Hong Wuxi	Semiconductor supply chain	N/A	N/A	\$0.7	N/A	N/A	N/A	\$0.7
InnoScience	Semiconductor chip design	N/A	\$0.2	\$0.4	N/A	\$0.1	N/A	\$0.7

Source: PitchBook • Geography: Global • As of September 30, 2025



DEAL AND EXIT ACTIVITY

Top 20 datacenter infrastructure companies by VC exit value (\$B)

Company	Segment	2020 exit value	2021 exit value	2022 exit value	2023 exit value	2024 exit value	2025 exit value	2020-2025 total exit value
CoreWeave	Operators & service providers	N/A	N/A	N/A	N/A	N/A	\$17.1	\$17.1
DigitalOcean	Operators & service providers	N/A	\$4.2	N/A	N/A	N/A	N/A	\$4.2
Kingsoft Cloud	Operators & service providers	\$3.0	N/A	N/A	N/A	N/A	N/A	\$3.0
Core Scientific	Facilities & power infrastructure	N/A	N/A	\$3.0	N/A	N/A	N/A	\$3.0
EdgeConneX	Facilities & power infrastructure	\$2.8	N/A	N/A	N/A	N/A	N/A	\$2.8
UCloud Technology	Operators & service providers	\$1.7	N/A	N/A	N/A	N/A	N/A	\$1.7
Cloud Light Technology	IT hardware & components	N/A	N/A	N/A	\$0.7	N/A	N/A	\$0.7
Epsagon	Software & orchestration	N/A	\$0.5	N/A	N/A	N/A	N/A	\$0.5
Volterra	Operators & service providers	N/A	\$0.4	N/A	N/A	N/A	N/A	\$0.4
QingCloud	Operators & service providers	N/A	\$0.3	N/A	N/A	N/A	N/A	\$0.3

Source: PitchBook • Geography: Global • As of September 30, 2025



DEAL AND EXIT ACTIVITY

Top 20 datacenter infrastructure companies by VC exit value (\$B) (continued)

Company	Segment	2020 exit value	2021 exit value	2022 exit value	2023 exit value	2024 exit value	2025 exit value	2020-2025 total exit value
Verne Global	Facilities & power infrastructure	N/A	\$0.3	N/A	N/A	N/A	N/A	\$0.3
Linktel Technologies	IT hardware & components	N/A	N/A	\$0.3	N/A	N/A	N/A	\$0.3
GRIID Infrastructure	Facilities & power infrastructure	N/A	N/A	N/A	\$0.3	N/A	N/A	\$0.3
Cloud Factory Technology	Operators & service providers	N/A	N/A	N/A	N/A	\$0.2	N/A	\$0.2
SAIHEAT	Facilities & power infrastructure	N/A	N/A	\$0.2	N/A	N/A	N/A	\$0.2
WarpStream	Software & orchestration	N/A	N/A	N/A	N/A	\$0.1	N/A	\$0.1
Block Mining	Facilities & power infrastructure	N/A	N/A	N/A	N/A	\$0.1	N/A	\$0.1
Bee Computing	Facilities & power infrastructure	N/A	\$0.1	N/A	N/A	N/A	N/A	\$0.1
Big Switch Networks	Software & orchestration	\$0.1	N/A	N/A	N/A	N/A	N/A	\$0.1
Yada Electronics	IT hardware & components	N/A	N/A	N/A	\$0.1	N/A	N/A	\$0.1

Source: PitchBook • Geography: Global • As of September 30, 2025



DEAL AND EXIT ACTIVITY

Top 20 semiconductor computing companies by VC exit value (\$B)

Company	Segment	2020 exit value	2021 exit value	2022 exit value	2023 exit value	2024 exit value	2025 exit value	2020-2025 total exit value
ASR Microelectronics	Semiconductor chip design	N/A	N/A	\$9.7	N/A	N/A	N/A	\$9.7
Horizon Robotics	Semiconductor chip design	N/A	N/A	N/A	N/A	\$6.0	N/A	\$6.0
GalaxyCore	Semiconductor chip design	N/A	\$5.0	N/A	N/A	N/A	N/A	\$5.0
Astera Labs	Semiconductor chip design	N/A	N/A	N/A	N/A	\$4.9	N/A	\$4.9
Nexchip Semiconductor	Semiconductor supply chain	N/A	N/A	N/A	\$4.3	N/A	N/A	\$4.3
SMEC	Semiconductor supply chain	N/A	N/A	N/A	\$4.2	N/A	N/A	\$4.2
Vanchip Technologies	Semiconductor chip design	N/A	N/A	\$3.8	N/A	N/A	N/A	\$3.8
ECARX	Semiconductor chip design	N/A	N/A	\$3.4	N/A	N/A	N/A	\$3.4
Cambricon Technologies	Semiconductor chip design	\$3.3	N/A	N/A	N/A	N/A	N/A	\$3.3
E-Town Semiconductor	Semiconductor supply chain	N/A	N/A	N/A	N/A	N/A	\$3.1	\$3.1

Source: PitchBook • Geography: Global • As of September 30, 2025



DEAL AND EXIT ACTIVITY

Top 20 semiconductor computing companies by VC exit value (\$B) (continued)

Company	Segment	2020 exit value	2021 exit value	2022 exit value	2023 exit value	2024 exit value	2025 exit value	2020-2025 total exit value
Novosense Microelectronics	Semiconductor chip design	N/A	N/A	\$2.7	N/A	N/A	N/A	\$2.7
VeriSilicon	Semiconductor supply chain	\$2.4	N/A	N/A	N/A	N/A	N/A	\$2.4
Bestechnic	Semiconductor chip design	\$2.2	N/A	N/A	N/A	N/A	N/A	\$2.2
Centec Networks	Semiconductor chip design	N/A	N/A	N/A	\$2.1	N/A	N/A	\$2.1
Southchip	Semiconductor chip design	N/A	N/A	N/A	\$2.1	N/A	N/A	\$2.1
Black Sesame Technologies	Semiconductor chip design	N/A	N/A	N/A	N/A	\$1.9	N/A	\$1.9
Halo Microelectronics	Semiconductor chip design	N/A	N/A	\$1.9	N/A	N/A	N/A	\$1.9
Primarius	Semiconductor supply chain	N/A	\$1.7	N/A	N/A	N/A	N/A	\$1.7
Intellifusion	Semiconductor chip design	N/A	N/A	N/A	\$1.7	N/A	N/A	\$1.7
Bluetrum Technology	Semiconductor chip design	N/A	N/A	\$1.6	N/A	N/A	N/A	\$1.6

Source: PitchBook • Geography: Global • As of September 30, 2025



DEAL AND EXIT ACTIVITY

Top 20 datacenter infrastructure companies by PE deal value (\$B)

Company	Segment	2020 deal value	2021 deal value	2022 deal value	2023 deal value	2024 deal value	2025 deal value	2020-2025 total deal value
Vantage Data Centers	Facilities & power infrastructure	\$2.5	\$1.5	N/A	\$2.7	\$9.2	\$2.2	\$18.1
AirTrunk	Facilities & power infrastructure	\$1.8	N/A	N/A	N/A	\$16.0	N/A	\$17.9
CyrusOne	Facilities & power infrastructure	N/A	\$15.0	N/A	N/A	N/A	N/A	\$15.0
Aligned Data Centers	Facilities & power infrastructure	N/A	N/A	N/A	\$0.0	\$0.0	\$12.0	\$12.1
Switch	Facilities & power infrastructure	N/A	N/A	\$11.0	\$0.0	N/A	N/A	\$11.0
CoreSite Realty	Facilities & power infrastructure	N/A	\$10.1	N/A	N/A	N/A	N/A	\$10.1
QTS Realty Trust	Facilities & power infrastructure	N/A	\$10.0	N/A	N/A	N/A	N/A	\$10.0
Compass Datacenters	Facilities & power infrastructure	N/A	N/A	N/A	\$5.7	N/A	N/A	\$5.7
DATA4	Facilities & power infrastructure	N/A	N/A	N/A	\$3.8	N/A	\$0.1	\$3.9
ChinData Group	Facilities & power infrastructure	N/A	N/A	N/A	\$3.2	N/A	N/A	\$3.2

Source: PitchBook • Geography: Global • As of September 30, 2025



DEAL AND EXIT ACTIVITY

Top 20 datacenter infrastructure companies by PE deal value (\$B) (continued)

Company	Segment	2020 deal value	2021 deal value	2022 deal value	2023 deal value	2024 deal value	2025 deal value	2020-2025 total deal value
EdgeCore Internet Real Estate	Facilities & power infrastructure	N/A	N/A	\$1.2	N/A	\$1.9	N/A	\$3.1
DataBank	Facilities & power infrastructure	N/A	N/A	N/A	\$0.2	\$2.9	N/A	\$3.1
Echelon Data Centres	Facilities & power infrastructure	\$1.1	N/A	N/A	N/A	\$1.9	N/A	\$3.0
EXA Infrastructure	Facilities & power infrastructure	\$2.8	N/A	N/A	N/A	N/A	N/A	\$2.8
EdgeConneX	Facilities & power infrastructure	\$2.8	N/A	N/A	N/A	N/A	N/A	\$2.8
Ensono	Operators & service providers	N/A	\$1.7	N/A	N/A	N/A	N/A	\$1.7
Couchbase	Software & orchestration	N/A	N/A	N/A	N/A	N/A	\$1.5	\$1.5
Nasuni	IT hardware & components	N/A	N/A	\$0.1	N/A	\$1.3	N/A	\$1.4
Zayo Group	Facilities & power infrastructure	\$1.4	N/A	N/A	N/A	N/A	N/A	\$1.4
eStruxture Data Centers	Facilities & power infrastructure	N/A	N/A	N/A	N/A	\$1.3	N/A	\$1.3

Source: PitchBook • Geography: Global • As of September 30, 2025



DEAL AND EXIT ACTIVITY

Top 20 semiconductor computing companies by PE deal value (\$B)

Company	Segment	2020 deal value	2021 deal value	2022 deal value	2023 deal value	2024 deal value	2025 deal value	2020-2025 total deal value
Tsinghua Unigroup	Semiconductor chip design	\$2.5	\$1.5	N/A	\$2.7	\$9.2	\$2.2	\$18.1
Shinko Electric Industries	Semiconductor supply chain	\$1.8	N/A	N/A	N/A	\$16.0	N/A	\$17.9
Altera	Semiconductor chip design	N/A	\$15.0	N/A	N/A	N/A	N/A	\$15.0
UNISOC	Semiconductor chip design	N/A	N/A	N/A	\$0.0	\$0.0	\$12.0	\$12.1
On-Bright Electronics	Semiconductor chip design	N/A	N/A	\$11.0	\$0.0	N/A	N/A	\$11.0
ASE Kunshan	Semiconductor supply chain	N/A	\$10.1	N/A	N/A	N/A	N/A	\$10.1
FICT	Semiconductor supply chain	N/A	\$10.0	N/A	N/A	N/A	N/A	\$10.0
Wuhan Xinxin Semiconductor	Semiconductor chip design	N/A	N/A	N/A	\$5.7	N/A	N/A	\$5.7
SJSemi	Semiconductor supply chain	N/A	N/A	N/A	\$3.8	N/A	\$0.1	\$3.9
KL-Tech	Semiconductor supply chain	N/A	N/A	N/A	\$3.2	N/A	N/A	\$3.2

Source: PitchBook • Geography: Global • As of September 30, 2025



DEAL AND EXIT ACTIVITY

Top 20 semiconductor computing companies by PE deal value (\$B) (continued)

Company	Segment	2020 deal value	2021 deal value	2022 deal value	2023 deal value	2024 deal value	2025 deal value	2020-2025 total deal value
SK Key Foundry	Semiconductor supply chain	N/A	N/A	\$1.2	N/A	\$1.9	N/A	\$3.1
Micross Components	Semiconductor chip design	N/A	N/A	N/A	\$0.2	\$2.9	N/A	\$3.1
Goermicro	Semiconductor supply chain	\$1.1	N/A	N/A	N/A	\$1.9	N/A	\$3.0
Analogix	Semiconductor chip design	\$2.8	N/A	N/A	N/A	N/A	N/A	\$2.8
Tianyue Advanced Technology	Semiconductor supply chain	\$2.8	N/A	N/A	N/A	N/A	N/A	\$2.8
SmartSens Technology	Semiconductor chip design	N/A	\$1.7	N/A	N/A	N/A	N/A	\$1.7
Changfei Advanced Semiconductor	Semiconductor supply chain	N/A	N/A	N/A	N/A	N/A	\$1.5	\$1.5
Epiworld	Semiconductor supply chain	N/A	N/A	\$0.1	N/A	\$1.3	N/A	\$1.4
Recif Technologies	Semiconductor supply chain	\$1.4	N/A	N/A	N/A	N/A	N/A	\$1.4
SiFusion	Semiconductor supply chain	N/A	N/A	N/A	N/A	\$1.3	N/A	\$1.3

Source: PitchBook • Geography: Global • As of September 30, 2025



DEAL AND EXIT ACTIVITY

Top 20 datacenter infrastructure companies by PE exit value (\$B)

Company	Segment	2020 exit value	2021 exit value	2022 exit value	2023 exit value	2024 exit value	2025 exit value	2020-2025 total exit value
AirTrunk	Facilities & power infrastructure	\$1.8	N/A	N/A	N/A	\$16.0	N/A	\$17.9
ChinData Group	Facilities & power infrastructure	\$4.3	N/A	N/A	N/A	N/A	\$3.9	\$8.2
Compass Datacenters	Facilities & power infrastructure	N/A	N/A	N/A	\$5.7	N/A	N/A	\$5.7
DATA4	Facilities & power infrastructure	N/A	N/A	N/A	\$3.8	N/A	N/A	\$3.8
OVH Groupe	Operators & service providers	N/A	\$3.7	N/A	N/A	N/A	N/A	\$3.7
Rackspace Technology	Operators & service providers	\$3.5	N/A	N/A	N/A	N/A	N/A	\$3.5
Ensono	Operators & service providers	N/A	\$1.7	N/A	N/A	N/A	N/A	\$1.7
Teraco Data Environments	Facilities & power infrastructure	N/A	N/A	\$1.7	N/A	N/A	N/A	\$1.7
eStruxture Data Centers	Facilities & power infrastructure	N/A	N/A	N/A	N/A	\$1.3	N/A	\$1.3
IREN	Facilities & power infrastructure	N/A	\$1.3	N/A	N/A	N/A	N/A	\$1.3

Source: PitchBook • Geography: Global • As of September 30, 2025



DEAL AND EXIT ACTIVITY

Top 20 datacenter infrastructure companies by PE exit value (\$B) (continued)

Company	Segment	2020 exit value	2021 exit value	2022 exit value	2023 exit value	2024 exit value	2025 exit value	2020-2025 total exit value
Cyxtera Technologies	Facilities & power infrastructure	N/A	\$1.3	N/A	N/A	N/A	N/A	\$1.3
EdgeCore Internet Real Estate	Facilities & power infrastructure	N/A	N/A	\$1.2	N/A	N/A	N/A	\$1.2
Nasuni	IT hardware & components	N/A	N/A	N/A	N/A	\$1.2	N/A	\$1.2
Green.ch	Facilities & power infrastructure	N/A	N/A	N/A	N/A	N/A	\$1.1	\$1.1
Nabiax	Facilities & power infrastructure	N/A	N/A	N/A	N/A	\$1.1	N/A	\$1.1
BigBear.ai Holdings	Software & orchestration	N/A	\$1.0	N/A	N/A	N/A	N/A	\$1.0
Green Mountain	Facilities & power infrastructure	N/A	\$0.9	N/A	N/A	N/A	N/A	\$0.9
Silent-Aire Manufacturing	Facilities & power infrastructure	N/A	\$0.8	N/A	N/A	N/A	N/A	\$0.8
HVR Software	Software & orchestration	N/A	\$0.7	N/A	N/A	N/A	N/A	\$0.7
Anord Mardix	Facilities & power infrastructure	N/A	\$0.5	N/A	N/A	N/A	N/A	\$0.5

Source: PitchBook • Geography: Global • As of September 30, 2025



DEAL AND EXIT ACTIVITY

Top 20 semiconductor computing companies by PE exit value (\$B)

Company	Segment	2020 exit value	2021 exit value	2022 exit value	2023 exit value	2024 exit value	2025 exit value	2020-2025 total exit value
Arm	Semiconductor supply chain	N/A	N/A	N/A	\$52.3	N/A	N/A	\$52.3
KIOXIA	Semiconductor chip design	N/A	N/A	N/A	N/A	\$4.9	N/A	\$4.9
Shenzhen Longsys Electronics	Semiconductor chip design	N/A	N/A	\$3.1	N/A	N/A	N/A	\$3.1
Allegro Microsystems	Semiconductor chip design	\$2.3	N/A	N/A	N/A	N/A	N/A	\$2.3
JOULWATT	Semiconductor chip design	N/A	N/A	\$2.1	N/A	N/A	N/A	\$2.1
Empyrean Technology	Semiconductor supply chain	N/A	N/A	\$2.1	N/A	N/A	N/A	\$2.1
Kokusai Electric	Semiconductor supply chain	N/A	N/A	N/A	\$2.0	N/A	N/A	\$2.0
Cmsemicon	Semiconductor chip design	N/A	N/A	\$1.8	N/A	N/A	N/A	\$1.8
Chipmore Technology	Semiconductor supply chain	N/A	N/A	N/A	\$1.7	N/A	N/A	\$1.7
SmartSens Technology	Semiconductor chip design	N/A	N/A	\$1.7	N/A	N/A	N/A	\$1.7

Source: PitchBook • Geography: Global • As of September 30, 2025



DEAL AND EXIT ACTIVITY

Top 20 semiconductor computing companies by PE exit value (\$B) (continued)

Company	Segment	2020 exit value	2021 exit value	2022 exit value	2023 exit value	2024 exit value	2025 exit value	2020-2025 total exit value
Orbbec	Semiconductor chip design	N/A	N/A	\$1.7	N/A	N/A	N/A	\$1.7
Ampleon	Semiconductor chip design	N/A	N/A	\$1.5	N/A	N/A	N/A	\$1.5
Semitronix	Semiconductor supply chain	N/A	N/A	\$1.3	N/A	N/A	N/A	\$1.3
Chengdu Sino-Microelectronics	Semiconductor chip design	N/A	N/A	N/A	N/A	\$1.2	N/A	\$1.2
C*Core Technology	Semiconductor supply chain	N/A	N/A	\$1.2	N/A	N/A	N/A	\$1.2
Jadard Tech	Semiconductor chip design	N/A	N/A	\$1.1	N/A	N/A	N/A	\$1.1
indie Semiconductor	Semiconductor chip design	\$1.1	N/A	N/A	N/A	N/A	N/A	\$1.1
National Silicon Industry Group	Semiconductor supply chain	\$1.0	N/A	N/A	N/A	N/A	N/A	\$1.0
Tianyue Advanced Technology	Semiconductor supply chain	N/A	\$0.9	N/A	N/A	N/A	N/A	\$0.9
Valens Semiconductor	Semiconductor chip design	N/A	\$0.9	N/A	N/A	N/A	N/A	\$0.9

Source: PitchBook • Geography: Global • As of September 30, 2025



Advanced computing PE and VC ecosystem market map

This market map is an overview of PE- and VC-backed companies that had received private capital or other notable private investments as of Q3 2025. Some categories have been combined to simplify the market map for readers. See the [taxonomy](#) for the full list of categories. [Click to view the full map on the PitchBook Platform.](#)

1 Semiconductor chip design

GPUs, CPUs, AI accelerators, & edge AI



Memory



Networking ICs



Analog & power ICs



Other chip design



Optoelectronics, sensors & discretes



2 Semiconductor supply chain

Advanced packaging



Capital equipment



EDA & IP



Foundry



Materials



OSAT



3 Datacenter infrastructure

Operators & service providers



Facilities & power infrastructure



Software & orchestration



IT hardware & components



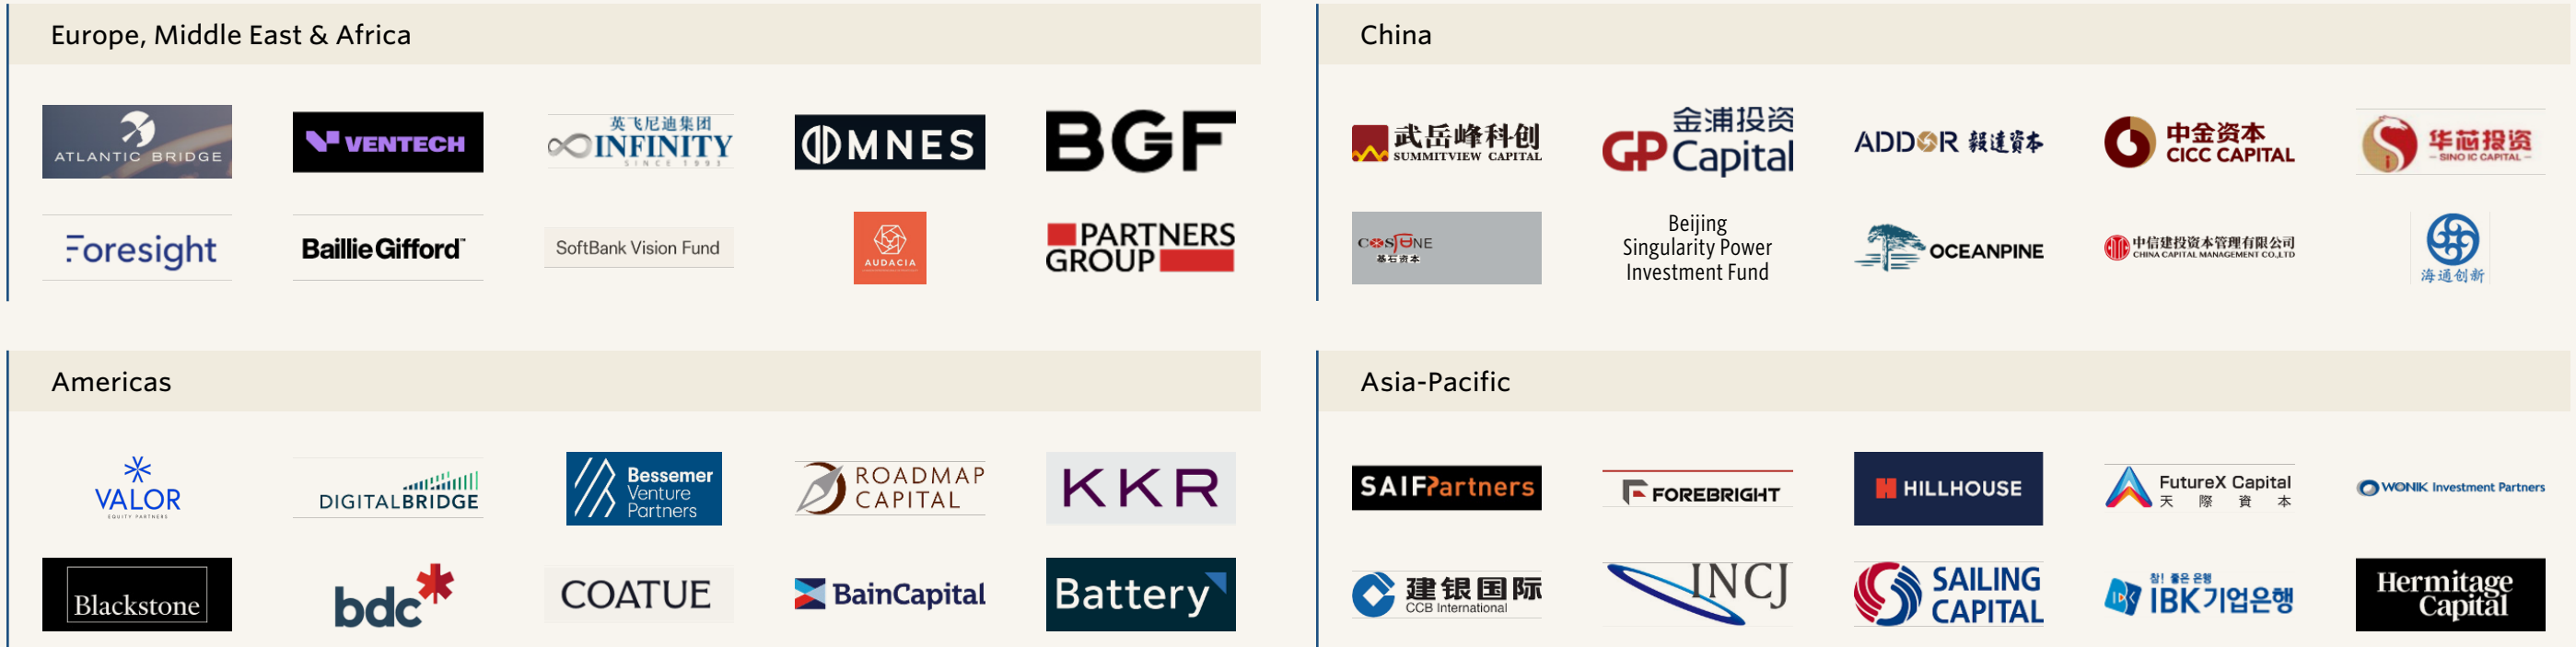
4 Quantum computing





Advanced computing PE investor map


This investor map is a representative overview of active investors in global buyouts and growth equity. Investors are ranked by their investment count in advanced computing from 2020 to Q3 2025.





Advanced computing VC investor map

This investor map is a representative overview of active investors in venture capital. Investors are ranked by their investment count in advanced computing from 2020 to Q3 2025.





Taxonomy

Semiconductor computing

Sector	Segment	Category	Description
Semiconductor	Semiconductor chip design	GPUs/CPUs	This subsegment combines primary processors for AI (GPUs) and general computing (CPUs). GPUs are the engines for AI factories. CPUs remain the workhorses for system control and traditional workloads.
Semiconductor	Semiconductor chip design	Memory (DRAM & NAND)	This subsegment includes IDMs that design and produce memory, including DRAM, NAND flash, and HBM. This is a commodity market essential for all systems. Pricing and profitability are highly cyclical.
Semiconductor	Semiconductor chip design	Networking & connectivity ICs	This subsegment includes chips that power network switches, routers, and data processing units. Key trends in this subsegment are increasing speed and offloading tasks from the CPU. These chips form the backbone of the datacenter network.
Semiconductor	Semiconductor chip design	Analog, mixed-signal & microcontroller units	This subsegment includes semiconductors that process real-world signals or perform dedicated control functions. They are used for power regulation, sensor monitoring, and system management. They are high-volume, lower-cost components.
Semiconductor	Semiconductor chip design	Power management ICs (PMICs)	This subsegment includes specialized chips responsible for managing, converting, and regulating power within a system. Efficient power management at the board level is critical for AI servers. These components are essential for maximizing power efficiency.
Semiconductor	Semiconductor chip design	Optoelectronics	This subsegment includes components that convert electrical signals to optical signals. This is foundational for high-speed fiber-optic communication inside datacenters. Growth in this subsegment is driven by the move to faster network speeds.
Semiconductor	Semiconductor chip design	Sensors, discretes & passives	This subsegment is a catch-all for foundational electronic components. This includes sensors for monitoring temperature, discrete transistors, and passive components. They are the fundamental building blocks of every electronic circuit board.
Semiconductor	Semiconductor chip design	AI accelerators (non-GPU)	This subsegment includes companies focused on creating specialized AI accelerator chips with architectures distinct from GPUs. These are often application-specific integrated circuits (ASICs) designed for specific AI workloads, such as inference. These chips include custom silicon from hyperscalers and startups.



TAXONOMY

Semiconductor computing (continued)

Sector	Segment	Category	Description
Semiconductor	Semiconductor chip design	Edge AI chips	This subsegment includes specialized, power-efficient processors designed to run AI inference on devices outside the datacenter. The end market is distinct, focused on Internet of Things (IoT), automotive, and mobile devices. The growth trend is parallel to the AI datacenter boom.
Semiconductor	Semiconductor chip design	Other semiconductor chip design	This catch-all subsegment for the semiconductor chip design segment captures companies whose primary focus is not covered by other subsegments. This includes designers of highly specialized ASICs for non-AI tasks and companies with highly diversified or legacy product lines where design is not the primary reported business. It ensures comprehensive coverage of the semiconductor chip design landscape.
Semiconductor	Semiconductor supply chain	Foundry	Foundries are the factories that physically fabricate silicon wafers. Their primary value is manufacturing-process leadership and scale. This is a highly capital-intensive business with immense barriers to entry.
Semiconductor	Semiconductor supply chain	EDA & engineering software	EDA provides the mission-critical software used to design and verify complex ICs. The software enables the entire chip-creation workflow.
Semiconductor	Semiconductor supply chain	IP	IP companies design and license prebuilt, verifiable circuit designs to chip designers. This accelerates time to market by providing proven technology blocks. This is a high-leverage, R&D-intensive business model.
Semiconductor	Semiconductor supply chain	Semiconductor materials & wafers	This subsegment supplies the foundational raw materials for chip production. Purity and quality are paramount. This is a concentrated market of materials science experts.
Semiconductor	Semiconductor supply chain	WFE	WFE companies build the complex machinery used inside foundries to manufacture semiconductors. These machines represent some of the most sophisticated technology ever developed. The market is an oligopoly with several monopoly submarkets.
Semiconductor	Semiconductor supply chain	OSAT	OSAT companies perform the "back-end" process of slicing wafers, packaging them into chips, and testing them. This is a critical, high-volume manufacturing step. The industry is consolidating around leaders with scale.



TAXONOMY

Semiconductor computing (continued)

Sector	Segment	Category	Description
Semiconductor	Semiconductor supply chain	Advanced packaging	This subsegment focuses on cutting-edge techniques for connecting multiple dies within a single package (chipselets). These techniques are crucial for high-performance AI accelerators. This is a critical bottleneck for AI chip production.
Semiconductor	Semiconductor supply chain	Other semiconductor supply chain	This catch-all subsegment for the semiconductor supply chain segment captures companies that support the manufacturing process but fall outside primary categories. This includes specialty materials suppliers (beyond raw wafers), metrology and inspection equipment, or third-party testing and analysis labs that are not full-scale OSAT providers. It ensures comprehensive coverage of the semiconductor supply chain landscape. This subsegment includes companies with highly diversified product lines.
Quantum computing	Quantum computing	Quantum hardware	This subsegment includes companies focused on building the physical quantum processing units (QPUs). Different physical modalities are being explored, including superconducting circuits and trapped ions. The primary challenge is increasing the number of stable, high-fidelity qubits.
Quantum computing	Quantum computing	Quantum software & cloud services	This subsegment focuses on building the software stack to program and utilize quantum computers. This includes developing quantum algorithms and offering platform access to various hardware back ends. The goal is to abstract the complexity of the underlying physics.



TAXONOMY

Datacenter infrastructure

Sector	Segment	Category	Description
Datacenter facilities	Datacenter facilities & power infrastructure	Physical site infrastructure	This subsegment includes the physical datacenter buildings and the firms that construct them. This is the foundation upon which all digital infrastructure is built. It includes operators that lease space and engineering firms.
Datacenter facilities	Datacenter facilities & power infrastructure	Power management & distribution	This subsegment provides equipment that manages and distributes high-voltage power. With rack power densities soaring due to AI, this equipment is foundational. This is the central nervous system of datacenter power.
Datacenter facilities	Datacenter facilities & power infrastructure	Thermal management (cooling)	This subsegment includes all technologies used to remove waste heat from IT equipment. The industry is rapidly shifting from air cooling to liquid cooling. This is a key enabling technology for AI factories.
Datacenter facilities	Datacenter facilities & power infrastructure	Backup power & generation	This subsegment provides on-site power generation and energy storage. These systems are necessary to keep a datacenter running during a grid outage. Reliability and runtime are the key metrics.
Datacenter facilities	Datacenter facilities & power infrastructure	Energy & utilities	This subsegment represents the power providers that supply primary power to datacenters. It includes traditional electric utilities. It also includes operators that arbitrage power costs for computation.
Datacenter facilities	Datacenter facilities & power infrastructure	Other datacenter facilities & power infrastructure	This catch-all subsegment for the datacenter facilities & power infrastructure segment captures companies providing niche physical infrastructure or services. This may include datacenter REITs, specialized monitoring and security services, or providers of modular or containerized datacenter solutions not covered elsewhere. It ensures comprehensive coverage of the datacenter facilities & power infrastructure landscape. This subsegment includes companies with highly diversified product lines.
Datacenter IT	Datacenter IT hardware & components	Computing infrastructure	This subsegment includes servers that perform core processing tasks. It covers general-purpose servers to high-density systems for AI. Differentiation lies in architecture, density, and management software.
Datacenter IT	Datacenter IT hardware & components	Storage infrastructure	This subsegment includes systems designed for storing and retrieving large amounts of data. Key metrics are input/output operations per second, throughput, and cost per gigabyte. The market is shifting from mechanical drives to all-flash systems.



TAXONOMY

Datacenter infrastructure (continued)

Sector	Segment	Category	Description
Datacenter IT	Datacenter IT hardware & components	Networking infrastructure	This subsegment provides the switches, routers, and interconnects that form the datacenter network fabric. The trend is toward higher speeds to support massive AI data flows. The network operating system is a key differentiator.
Datacenter IT	Datacenter IT hardware & components	Specialized IT components	This subsegment includes other critical hardware components that are not full systems. This includes power supplies and cabling. While often overlooked, they are a significant part of the bill of materials.
Datacenter IT	Datacenter IT hardware & components	Other datacenter IT hardware & components	This catch-all subsegment for the datacenter IT hardware & components segment captures companies providing specialized hardware. This includes manufacturers of niche server components, custom-built "white box" server integrators, or providers of specialized testing and validation hardware for IT infrastructure. It ensures comprehensive coverage of the datacenter IT hardware & components landscape. This subsegment includes companies with highly diversified product lines.
Datacenter operators	Datacenter operators & service providers	Edge & content delivery	This subsegment includes companies that operate globally distributed networks of smaller edge datacenters. They deliver content and applications with low latency. This is critical for real-time applications and security.
Datacenter operators	Datacenter operators & service providers	Cloud platforms & infrastructure as a service (IaaS)/platform as a service (PaaS)	This subsegment is defined by providers that offer foundational IaaS and PaaS. They provide raw computing and storage on a pay-as-you-go basis. They are the primary drivers of datacenter hardware consumption.
Datacenter operators	Datacenter operators & service providers	Managed services & interconnection	This subsegment includes services that help enterprises manage hybrid and multicloud environments. It also includes the vital interconnection services that link networks and clouds together. These are the "glue" that holds the hybrid cloud world together.
Datacenter operators	Datacenter operators & service providers	Other datacenter operators & service providers	This catch-all subsegment for the datacenter operators & service providers segment captures service providers that do not fit into the primary subsegments. These may include specialized "bare metal" infrastructure providers, datacenter interconnection specialists (beyond managed services), or platform-agnostic cloud brokerage and management firms. This subsegment ensures comprehensive coverage of the datacenter operators & service providers landscape. It includes companies with highly diversified product lines.



TAXONOMY

Datacenter infrastructure (continued)

Sector	Segment	Category	Description
Datacenter IT	Datacenter software & orchestration	Virtualization & containerization	This subsegment includes companies that develop software that abstracts the underlying physical hardware. It allows multiple operating systems or applications to run on a single server. Kubernetes is the de facto standard for container orchestration.
Datacenter IT	Datacenter software & orchestration	Data platforms & analytics	This subsegment provides software for storing, processing, and analyzing large datasets. These platforms are essential for turning raw data into business intelligence. The shift to the cloud has transformed this category.
Datacenter IT	Datacenter software & orchestration	AI & ML development platforms	This subsegment includes the software platforms and foundation models used to build, train, and deploy AI applications. This is the core software enabling the current AI boom. It represents a significant new investment area.
Datacenter IT	Datacenter software & orchestration	Observability & management software	This subsegment includes companies whose software provides monitoring and insights into the performance and health of applications and infrastructure. Such software combines logs, metrics, and traces to help operators detect and resolve issues quickly. It is critical for maintaining uptime in complex systems.
Datacenter IT	Datacenter software & orchestration	Other datacenter software & orchestration	This catch-all subsegment for the datacenter software & orchestration segment captures software companies with tools not covered in other categories. This includes niche workload automation, IT service management tools adapted for datacenters, and specialized database management and optimization software. This subsegment ensures comprehensive coverage of the datacenter software & orchestration landscape. It includes companies with highly diversified product lines.



Glossary

Term	Description	Segment	Tier
Adaptive computing	Hardware (such as field-programmable gate arrays) that can be reconfigured after manufacturing. Adaptive computing provides infrastructure future-proofing, allowing hardware to adapt to rapidly changing AI algorithms without replacement.	Semiconductor chip design	Chips & hardware
Advanced packaging	Manufacturing techniques that combine multiple distinct dies into a single package. Advanced packaging extends Moore's Law, allowing chipmakers to continue increasing performance and revenue even as individual transistor scaling slows down.	Semiconductor supply chain	Chips & hardware
Agentic AI	AI systems capable of autonomous action and decision-making. Agentic AI unlocks high-value automation markets requiring significantly higher reliability and security than simple chatbots.	Datacenter software & orchestration	Foundation models
AI accelerator	A specialized processor, such as a GPU, ASIC, or tensor processing unit (TPU), optimized for parallel processing tasks. AI accelerators are the core revenue driver of the sector, and demand for these chips is the primary indicator of the entire industry's growth.	Datacenter IT hardware & components	Chips & hardware
AI cluster	A group of connected servers working together as a single system. An AI cluster is the operational unit of AI, and understanding cluster size helps investors gauge the scale of training capabilities.	Datacenter IT hardware & components	Chips & hardware
AI factories	Specialized datacenters designed specifically for the massive scale, power, and thermal requirements of AI training, representing a new dedicated asset class commanding premium valuations distinct from general enterprise datacenters.	Datacenter facilities & power infrastructure	Data center operations
AI inference	The process of using a trained model to make predictions or generate content. This revenue-generating phase of AI is distinct from training, where ongoing value is delivered to users.	Datacenter software & orchestration	AI-native applications & services
AI native	Applications or companies built from the ground up with AI as the core value proposition. AI-native entities represent the highest disruption potential as they reimagine workflows entirely.	Datacenter software & orchestration	AI-native applications & services
AI networking	High-bandwidth communication infrastructure designed specifically to connect accelerators within a cluster. AI networking acts as the performance backbone allowing thousands of GPUs to function as a single computer, enabling massive model training.	Datacenter IT hardware & components	Chips & hardware



GLOSSARY

Term	Description	Segment	Tier
Air cooling	The traditional method of cooling servers using fans and air exchange. Air cooling's thermal limit (approximately 30 kW/rack) defines the technological obsolescence boundary, as facilities relying solely on this cannot host modern AI clusters.	Datacenter facilities & power infrastructure	Energy, power, cooling & real estate
AI reasoning	The capability of an AI to simulate complex, multistep logical deductions (System 2 thinking). AI reasoning drives a massive increase in inference-time computing demand, acting as a revenue multiplier.	Datacenter software & orchestration	Foundation models
AI storage	Storage systems optimized for high-throughput, parallel data access to feed GPUs. They are essential for preventing input/output starvation, ensuring expensive GPUs are not sitting idle waiting for data during training.	Datacenter IT hardware & components	Chips & hardware
AI technology stack	The layered framework of hardware and software needed to deploy and manage AI. Understanding its components is essential for identifying bottlenecks and market share distribution across the value chain.	Datacenter software & orchestration	Software infrastructure
AI workloads	The specific computational tasks (such as training, fine-tuning, and inference) performed by AI. Different workloads require different hardware, so understanding the mix dictates infrastructure demand.	Datacenter software & orchestration	Software infrastructure
ARC-AGI-1	The original benchmark for abstract reasoning. ARC-AGI-1 was a historical baseline whose saturation by models drove the need for harder tests, illustrating the rapid pace of AI advancement.	Datacenter software & orchestration	Foundation models
ARC-AGI-2	The 2024 successor to the ARC benchmark designed to be unsolvable by pattern matching. It serves as the gold-standard test for measuring genuine fluid intelligence and filtering out hype.	Datacenter software & orchestration	Foundation models
Artificial general intelligence (AGI)	An AI capable of understanding and learning any intellectual task a human can. AGI is the infinite demand driver whose achievement would fundamentally disrupt the global economy and maximize computing consumption.	Datacenter software & orchestration	Foundation models
Busway	A modular, overhead track system for distributing power to server racks. Busways provide the deployment speed and flexibility to allow operators to rapidly reconfigure power delivery for evolving, power-hungry AI hardware.	Datacenter facilities & power infrastructure	Energy, power, cooling & real estate



GLOSSARY

Term	Description	Segment	Tier
Chillers	Industrial mechanical systems that remove heat from a building's liquid-cooling loop via a refrigeration cycle. Their efficiency determines the overall power usage effectiveness and operational expenditure of a facility, directly influencing the long-term profitability of AI hosting.	Datacenter facilities & power infrastructure	Energy, power, cooling & real estate
Chiplet	A modular integrated circuit block designed to be combined with others. This architecture enables cost-effective customization and higher yields, allowing manufacturers to build specialized AI processors without the cost of monolithic dies.	Semiconductor chip design	Chips & hardware
Chip-on-Wafer-on-Substrate (CoWoS)	An advanced 3D packaging technology connecting logic (GPU) and memory (HBM) on a silicon interposer. It is the primary manufacturing bottleneck for NVIDIA/AMD high-end chips, directly dictating global AI accelerator supply.	Semiconductor supply chain	Chips & hardware
Cloud server provider (CSP)	A general term for companies providing cloud computing services. Distinguishing between Tier 1 (hyperscalers) and Tier 2 (neoclouds) is key to understanding competitive dynamics in the service layer.	Datacenter operators & service providers	Data center operations
Colocation	A service renting space, power, and cooling for customer-owned hardware. Colocation is essential for hybrid AI strategies, allowing enterprises to own their sensitive AI hardware while leveraging professional infrastructure.	Datacenter operators & service providers	Data center operations
Computer room air handlers (CRAHs)	Large air conditioning units that circulate cool air through a data hall. Their inability to handle modern GPU heat loads signals the imminent capital expenditure shift away from legacy air cooling toward liquid-cooling solutions.	Datacenter facilities & power infrastructure	Energy, power, cooling & real estate
Context window	The amount of information (tokens) a model can hold in short-term memory. Larger windows allow for analyzing entire books or codebases, a critical feature for enterprise utility and retrieval-augmented generation (RAG) applications.	Datacenter software & orchestration	Foundation models
Cooling towers	Large outdoor heat exchangers that cool water by evaporating a small portion of it into the atmosphere. Cooling towers define the environmental and regulatory limits (water usage) of an AI factory, often becoming a barrier to site approval.	Datacenter facilities & power infrastructure	Energy, power, cooling & real estate
Crystallized intelligence	The ability to use learned knowledge and experience (memorization). Crystallized intelligence represents the baseline capability of current LLMs, distinguishing commodity knowledge retrieval from high-value reasoning.	Datacenter software & orchestration	Foundation models



GLOSSARY

Term	Description	Segment	Tier
CUDA	NVIDIA's proprietary parallel computing platform and programming model. CUDA is the dominant software moat locking developers into NVIDIA hardware, allowing a massive competitive advantage via ecosystem lock-in.	Datacenter software & orchestration	Software infrastructure
Custom ASIC/ custom silicon	Chips designed for a single specific workload (such as inference) rather than general purposes. This is a strategic lever for hyperscalers to reduce capital expenditure and dependency on merchant vendors for massive-scale workloads.	Semiconductor chip design	Chips & hardware
Datacenter infrastructure management (DCIM)	Software for monitoring and managing physical datacenter assets. DCIM is the operational intelligence layer required to safely run AI factories at extreme densities without thermal failure.	Datacenter software & orchestration	Software infrastructure
Data lakehouse	A data architecture combining the flexibility of a data lake with the management of a data warehouse. Data lakehouses are the central data repository providing the governed, massive datasets required for training foundation models.	Datacenter software & orchestration	Software infrastructure
Data processing unit (DPU)	A specialized chip dedicated to handling datacenter infrastructure tasks (such as networking and security). A DPU offloads the CPU to focus on computing, significantly increasing the efficiency and security of the AI cloud.	Semiconductor chip design	Chips & hardware
Data warehouse	A centralized repository for storing and analyzing structured data. Data warehouses remain critical for the structured business data used to fine-tune enterprise AI models for specific corporate tasks.	Datacenter software & orchestration	Software infrastructure
Deep learning	A subset of ML based on artificial neural networks with representation learning. Deep learning is the scientific breakthrough enabling modern AI, dictating the entire hardware architecture of the sector.	Datacenter software & orchestration	AI-native applications & services
Digital twin	A virtual replica of a physical system used for simulation. Digital twins enable predictive maintenance and optimization of complex AI factories, reducing downtime and improving efficiency.	Datacenter software & orchestration	Software infrastructure
Edge AI	AI processing performed on local devices rather than the cloud. Edge AI unlocks low-latency revenue streams and drives demand for specialized, power-efficient inference chips.	Datacenter IT hardware & components	Chips & hardware



GLOSSARY

Term	Description	Segment	Tier
Embedded edge AI	The deployment of lightweight AI models directly onto low-power edge devices. This enables real-time decision-making without cloud connectivity, driving the next wave of IoT and smart device proliferation.	Datacenter operators & service providers	AI-native applications & services
Embedded silicon	Specialized processors (systems on chips/MCUs) integrated directly into physical devices to perform real-time AI inference locally. Embedded silicon enables the autonomy of physical AI by removing cloud latency and driving demand for high-efficiency edge chips.	Semiconductor chip design	Chips & hardware
Enterprise AI	The application of AI to optimize white-collar business processes (such as HR and finance). Enterprise AI represents a massive productivity market driving high-margin software revenue and stable inference demand.	Datacenter software & orchestration	AI-native applications & services
Extreme ultraviolet (EUV) lithography	A manufacturing process using extreme ultraviolet light to print sub-7-nm transistors. This is the ultimate geopolitical choke point, as only ASML machines can produce the advanced chips required for modern AI.	Semiconductor supply chain	Chips & hardware
Fast mode	The standard, low-latency inference path where a model generates tokens immediately without hidden reasoning steps. Fast mode is essential for real-time AI applications such as voice and chatbots.	Datacenter software & orchestration	Foundation models
Fluid intelligence	The ability to reason, solve novel problems, and identify patterns in new situations independent of training data. Achieving this is the primary hurdle and ultimate goal for AGI.	Datacenter software & orchestration	Foundation models
Foundry	A facility that manufactures chips for other companies. Foundries are the ultimate production gatekeeper, as TSMC's capacity effectively caps the maximum growth rate of the entire AI sector.	Semiconductor supply chain	Chips & hardware
Generative AI (GenAI)	AI capable of generating text, images, or other media in response to prompts. GenAI is the market-defining capability driving the current investment boom, shifting value from analysis to creation.	Datacenter software & orchestration	AI-native applications & services
GPU Direct RDMA	A technology allowing GPUs to access network data directly by bypassing the CPU. This is essential for maximizing training speed by removing the CPU bottleneck in massive, distributed AI training runs.	Datacenter IT hardware & components	Chips & hardware



GLOSSARY

Term	Description	Segment	Tier
Graduate-Level Google-Proof Q&A Benchmark (GPQA)	A dataset of Ph.D.-level science questions that benchmarks an AI's ability to perform expert-level reasoning, serving as a key metric for models targeting scientific R&D markets.	Datacenter software & orchestration	Foundation models
Graduate-Level Google-Proof Q&A Benchmark (GPQA) Diamond	The highest-difficulty subset of the GPQA benchmark containing Ph.D.-level scientific questions. It is the ultimate litmus test for super-expert reasoning, distinguishing frontier models from commodity models.	Datacenter software & orchestration	Foundation models
Gray space	The area of a datacenter dedicated to back-end infrastructure (such as power, cooling, and switchgear) rather than IT equipment. AI workloads require a significantly larger ratio of gray space to white space, increasing construction cost and complexity.	Datacenter facilities & power infrastructure	Energy, power, cooling & real estate
Hardware/software co-design	The engineering philosophy where chips and software libraries (such as CUDA) are designed simultaneously. This creates a high-performance moat that makes it nearly impossible for rivals to match performance with a siloed approach.	Semiconductor chip design	Chips & hardware
Hidden tokens	Intermediate reasoning steps generated by a model but not shown to the user. Hidden tokens are essential for complex problem solving; they increase computing consumption per query, boosting inference revenue.	Datacenter software & orchestration	Foundation models
High-bandwidth memory (HBM)	A type of memory architecture where 3D-stacked DRAM is placed directly next to the processor. HBM is the single largest hardware bottleneck for AI, as model speed is currently limited by memory bandwidth.	Semiconductor chip design	Chips & hardware
High-performance Ethernet	Standardized Ethernet evolved for low-latency, lossless transmission. It provides a commodity alternative to InfiniBand, driving broader, lower-cost adoption of AI networking across enterprises.	Datacenter IT hardware & components	Chips & hardware
Humanity's Last Exam (HLE)	A multidisciplinary benchmark of expert-level questions designed to be the final test of academic knowledge before AI surpasses human capability, marking a major valuation inflection point.	Datacenter software & orchestration	Foundation models
Hyperscale cloud platform	The massive cloud infrastructure services offered by major tech companies. They are the primary distribution channel for AI services and the largest purchasers of AI hardware.	Datacenter operators & service providers	Software infrastructure



GLOSSARY

Term	Description	Segment	Tier
Hypervisor	Software that creates and runs virtual machines. Hypervisors are a foundational technology for cloud computing, key to maximizing the utilization of expensive server hardware.	Datacenter software & orchestration	Software infrastructure
Industrial AI	The application of AI to physical industrial assets and processes (manufacturing). Industrial AI drives efficiency in the physical economy, creating demand for ruggedized hardware and distinct models.	Datacenter operators & service providers	AI-native applications & services
InfiniBand	A high-speed networking standard offering high throughput and low latency. InfiniBand defines the performance ceiling for AI training clusters, and its proprietary nature drives high margins for leaders like NVIDIA.	Datacenter IT hardware & components	Chips & hardware
Language model fidelity (LMF)	The measure of a model's strict adherence to factual ground truth and user intent. LMF is the fundamental safety requirement for deploying autonomous agentic AI in enterprise environments.	Datacenter software & orchestration	Foundation models
Large language model (LLM)	An AI model trained on vast amounts of text data to understand and generate human language. LLMs are the primary software assets driving current infrastructure investment and the competitive landscape.	Datacenter software & orchestration	Foundation models
Leading-edge process technologies	The most advanced semiconductor nodes (such as 3 nm and 2 nm) in production. Access to these nodes dictates chip performance and efficiency, serving as the primary competitive advantage for chip designers.	Semiconductor supply chain	Chips & hardware
Liquid cooling	Advanced thermal management utilizing fluid to absorb heat directly from components. Liquid cooling enables next-generation AI servers (over 40 kW/rack) and sustains computing scaling, acting as a major capital expenditure catalyst for the cooling market.	Datacenter facilities & power infrastructure	Energy, power, cooling & real estate
Machine learning operations (MLOps)	Tools and practices for deploying and maintaining AI models. MLOps is the operational scaffolding that transforms experimental models into reliable, revenue-generating enterprise products.	Datacenter software & orchestration	Software infrastructure
Memory wall	The disparity between the speed of the processor and the speed of memory data transfer. The memory wall is the fundamental architectural problem driving massive investment in HBM and interconnects, shaping the entire hardware road map.	Semiconductor chip design	Chips & hardware
Merchant silicon	Standardized, high-performance chips sold to any buyer (such as by Broadcom or AMD). This strategy promotes open standards and breaks proprietary vendor lock-in in the networking and computing layers.	Semiconductor chip design	Chips & hardware



GLOSSARY

Term	Description	Segment	Tier
Metrology	The science of measurement in chip manufacturing. Metrology is essential for yield management in advanced nodes, without which the cost of producing AI chips becomes prohibitive.	Semiconductor supply chain	Chips & hardware
Mixture of experts (MoE)	A model architecture that activates only a subset of parameters per token. MoE is key to economic viability by allowing models to scale to trillions of parameters while keeping inference costs manageable.	Datacenter software & orchestration	Foundation models
Model expert intelligence	The level of domain-specific expertise an LLM demonstrates on benchmarks. Model expert intelligence quantifies an AI's value for vertical-specific applications, driving enterprise adoption in those sectors.	Datacenter software & orchestration	Foundation models
Multimodal models	AI models capable of processing and generating multiple data types (such as text, images, and audio) simultaneously. Multimodal models are a massive computing multiplier requiring significantly more data throughput than text alone.	Datacenter software & orchestration	Foundation models
Neocloud	Specialized cloud providers focused solely on high-performance AI computing. Neoclouds offer a competitive alternative to hyperscalers, often providing newer hardware faster or at a lower cost for specific AI needs.	Datacenter operators & service providers	Data center operations
Neural networks	Computing systems loosely inspired by the biological neural networks of animal brains. Their matrix-math operations determine the hardware requirements (GPUs/TPUs) for the entire industry.	Datacenter software & orchestration	AI-native applications & services
Neuromorphic computing	Chip architectures modeled on the human brain's neural structure. Neuromorphic computing promises to slash energy consumption by orders of magnitude for inference, potentially revolutionizing edge AI economics.	Semiconductor chip design	Chips & hardware
Operational technology (OT)	Hardware and software that detects or controls physical devices. OT is the integration target for industrial AI, representing the legacy install base that must be upgraded or bridged for AI adoption.	Datacenter operators & service providers	AI-native applications & services
Original design manufacturers (ODMs)	Companies that design and manufacture products that are rebranded by others. ODMs are the hidden manufacturing engine behind the cloud, building the vast majority of the physical infrastructure for hyperscalers.	Datacenter IT hardware & components	Data center operations



GLOSSARY

Term	Description	Segment	Tier
Original equipment manufacturers (OEMs)	Companies that build servers using chips from vendors such as NVIDIA. They are the primary sales channel for reaching the massive enterprise market, distinct from the direct-to-cloud hyperscaler channel.	Datacenter operators & service providers	Data center operations
Outsourced semiconductor assembly & test (OSAT)	Companies that package and test manufactured chips. OSAT is a critical, often overlooked bottleneck where advanced packaging capacity determines the final supply volume of AI accelerators.	Semiconductor supply chain	Chips & hardware
Perfect fidelity	The theoretical state where an AI model's output is 100% accurate to the source material and user intent. This is the ultimate quality standard required for high-stakes sectors like healthcare.	Datacenter software & orchestration	Foundation models
Performance per watt (PPW)	A measure of computational output (Tokens/FLOPS) per unit of energy consumed. PPW is the core design philosophy for all modern hardware, directly determining the operational expenditure profitability of deploying AI at scale.	Datacenter IT hardware & components	Chips & hardware
Photonic interconnect	Data transmission technology using light instead of electricity between chips. This is the successor to copper interconnects necessary to solve the energy and bandwidth limits of future massive-scale AI clusters.	Datacenter IT hardware & components	Chips & hardware
Physical AI	AI systems that interact with the physical world through robotics or sensors. Physical AI is the bridge between digital intelligence and physical action, representing the next massive TAM expansion.	Datacenter operators & service providers	AI-native applications & services
Power density	The amount of electrical power (kW) delivered to a specific area or server rack. This defines the physical limit of an AI factory and acts as the primary constraint on revenue growth per square foot.	Datacenter facilities & power infrastructure	Energy, power, cooling & real estate
Power distribution unit (PDU)	An industrial-grade power strip that distributes and monitors electricity at the rack level. PDUs are the final point of control for optimizing power efficiency and allocating costs in high-density AI environments.	Datacenter IT hardware & components	Energy, power, cooling & real estate
Power quality	The consistency and stability of the voltage and amperage delivered to IT equipment. High quality is essential for hardware reliability, as fluctuations can cause instant failure of sensitive, multimillion-dollar AI accelerator clusters.	Datacenter facilities & power infrastructure	Energy, power, cooling & real estate



GLOSSARY

Term	Description	Segment	Tier
Pro mode	A service tier allocating premium computational resources to a user request. Scaling from low to high increases the context window and model size, serving as the primary lever for monetizing power users.	Datacenter software & orchestration	Foundation models
Prompt	The user-provided instruction or input given to a GenAI model. The prompt is the control surface for GenAI where complexity directly correlates to the complexity of the computing workload.	Datacenter software & orchestration	AI-native applications & services
Prompt engineering	The practice of optimizing inputs to an AI model to generate desired outputs. Prompt engineering is specialized human skill layer that maximizes model utility, bridging raw capability and business value.	Datacenter software & orchestration	AI-native applications & services
Quantum bit (qubit)	The fundamental unit of information in a quantum computer capable of representing multiple states simultaneously. Its stability dictates the commercial timeline for quantum computing as the successor to silicon acceleration.	Quantum computing	Chips & hardware
Rack density	The measure of computing power (number of GPUs/CPUs) packed into a single server rack. Maximizing this metric is key to capital expenditure efficiency for hyperscalers by minimizing the facility footprint required to host high-margin AI services.	Datacenter IT hardware & components	Energy, power, cooling & real estate
Rack-scale computing	An architecture where the entire server rack (power, cooling, and networking) is designed as a single, integrated computer to maximize component density and energy efficiency, driving next-generation capital expenditure design standards.	Datacenter IT hardware & components	Energy, power, cooling & real estate
Reasoning effort	A user-configurable parameter (low, medium, high) that dictates the depth of an AI's internal chain of thought. Higher settings exponentially increase computing cost and accuracy.	Datacenter software & orchestration	Foundation models
Retrieval-augmented generation (RAG)	A technique that enhances model accuracy by fetching external data to include in the prompt. RAG solves hallucination problems and is a critical enabler for enterprise adoption.	Datacenter software & orchestration	AI-native applications & services
ROCm	AMD's open-source software stack for GPU computing and the primary challenger to CUDA, whose maturity is the single biggest factor in determining if AMD can gain significant market share.	Datacenter software & orchestration	Software infrastructure



GLOSSARY

Term	Description	Segment	Tier
Semiconductor supply chain	The network of materials, equipment, and services required to make chips. The semiconductor supply chain highlights systemic risks where disruption in one link (such as neon gas or substrates) can halt the entire AI hardware industry.	Semiconductor supply chain	Chips & hardware
Semicustom silicon	A chip design model where a vendor integrates proprietary IP with a customer's specific requirements. This allows chipmakers to capture revenue from hyperscalers that want differentiated hardware without full custom design risk.	Semiconductor chip design	Chips & hardware
Small language model (SLM)	Lightweight models optimized for edge devices and cost-efficiency. SLMs are critical for scaling AI to mass-market devices (such as phones and laptops) where cloud latency and cost are prohibitive.	Datacenter software & orchestration	Foundation models
Sovereign AI	National policies mandating domestic AI development and data processing. Sovereign AI is a geopolitical driver creating guaranteed, government-backed markets for hardware and cloud providers.	Datacenter operators & service providers	Foundation models
Superintelligence	A hypothetical AI system with intellect far surpassing human cognitive abilities across all domains. Superintelligence frames the ultimate TAM and existential risk profile of the sector, driving long-term valuation.	Quantum computing	Foundation models
SWE-Bench	A benchmark tasking agents with resolving real GitHub issues. SWE-Bench is the industry standard for evaluating an AI's ability to write reliable code, directly correlating to its utility in software automation.	Datacenter software & orchestration	Foundation models
Switchgear	High-voltage equipment used to control and protect electrical equipment. Switchgear is currently a major supply-side bottleneck with long lead times that can delay AI factory construction and revenue generation by months.	Datacenter facilities & power infrastructure	Energy, power, cooling & real estate
Synthetic data	Artificial data generated by AI to train other AI. Synthetic data is the only scalable solution to the "data wall" (running out of human data), creating a new market for data generation services.	Datacenter software & orchestration	Foundation models
Thinking mode	An inference paradigm where the AI generates hidden chains of thought before responding; the settings (low/medium/high) dictate the amount of test-time computing and revenue per query.	Datacenter software & orchestration	Foundation models



GLOSSARY

Term	Description	Segment	Tier
Tokens	The basic units of text (words/characters) processed by an LLM. Tokens are the fundamental economic unit of AI used for consumption, pricing, and capacity planning calculations.	Datacenter software & orchestration	Foundation models
Tokens per dollar per watt	An efficiency metric measuring computational output relative to cost and energy. This is the ultimate comparative metric for AI platforms, quantifying the economic and energy efficiency of generating AI output.	Datacenter software & orchestration	Software infrastructure
Total cost of ownership (TCO)	The comprehensive metric (operational and capital expenditure) used to evaluate investments. Competitors use TCO to challenge premium pricing by demonstrating long-term cost-effectiveness of alternative stacks.	Datacenter software & orchestration	Software infrastructure
Transformer architecture	The deep learning architecture based on self-attention mechanisms. Transformer architecture is the foundational blueprint of modern AI essential to understanding the parallel computing and memory demands driving hardware capital expenditure.	Datacenter software & orchestration	Foundation models
Vision AI	AI designed to interpret and analyze visual information from images or video. Visual AI is a key driver for edge AI hardware, enabling high-growth applications such as autonomous vehicles and security.	Datacenter operators & service providers	AI-native applications & services
World models	AI models that learn a physical simulation of the world rather than just language patterns. World models are the foundational software requirement for physical AI and robotics to function safely in reality.	Datacenter software & orchestration	Foundation models
Zero trust platforms	Integrated security suites that enforce the zero trust philosophy. They provide the operational security guarantee necessary for enterprises to deploy autonomous AI agents with confidence.	Datacenter software & orchestration	Software infrastructure



Methodology

The following search criteria was used on the PitchBook Platform to generate initial deal flow data for the advanced computing vertical.³

Industry

Industries, verticals & keywords: Semiconductors (industry), systems and information management (industry), data center, datacenters, quantum computing (emerging space)

Deal criteria

- Deal size: At least \$1 million
- Deal date: After January 1, 2020
- Size and date apply to the full transaction.

Deal types

Venture capital

- All VC stages
- All round numbers
- All series

Corporate/strategic M&A

- M&A/control transactions

Private equity

- All buyout types
- Other private equity types
 - Growth/expansion

IPO/liquidity

- Public investments
 - IPO


³: For this launch report, we excluded the quantum computing emerging space category in our search criteria to focus on semiconductor-related quantum deals. We will include quantum computing in our search criteria for future reports.

PitchBook Data, Inc.

Nizar Tarhuni Executive Vice President of Research and Market Intelligence
Paul Condra Global Head of Private Markets Research

Additional research

Industry research



Q4 2025 Analyst Note: Sovereign AI: The Trillion-Dollar Frontier

Download the report [here](#)



Q3 2025 AI VC Trends

Download the report [here](#)



Q3 2025 AI Public Comp Sheet and Valuation Guide

Download the report [here](#)



Q3 2025 Enterprise SaaS VC Trends

Download the report [here](#)

More research available at pitchbook.com/news/reports

PitchBook, a Morningstar company

©2025 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.