Bloomberg Intelligence BI

# Generative AI

## 2025 Outlook

## Gauging the Opportunities After DeepSeek Disrupts a Market Racing for $2 Trillion

The adoption of generative artificial intelligence and large language models (LLM) is already coursing through every segment of the technology sector and evolving rapidly, as incumbents and new entrants reimagine existing end markets to potentially generate about $1.8 trillion in revenue by 2032.

- **Focus Shifts from Pre-Training to Reasoning Models:** The pivot toward reasoning models powered by chain of thought and reinforcement learning will likely enable broader use of LLMs beyond text-based searches to analyzing images, audio and videos across a variety of form factors. Among the foundational model companies, OpenAI-Microsoft, Google, Meta, DeepSeek, Anthropic and Mistral continue to release new LLMs with chain-of-thought reasoning at inference time.
- **Gen AI Agents Set to Explode:** In addition to the established use cases around contract review and customer-service chatbots powered by LLMs, the integration of writing and coding assistants, and image and video creation tools using text and voice prompts will power gen AI agent deployments for consumer and enterprise uses to around $214 billion in spending by 2032.
- **AI Inference Model Efficiency After DeepSeek:** Most LLM companies are focused on driving model efficiency for achieving scale in inferencing. Instead of processing everything on the cloud, on-device AI adoption will be driven by a focus on privacy for features like writing text, translating language, generating images and videos with text-based prompts, and voice assistants.

**Featured in This Report:** This is the third update of Bloomberg Intelligence's deep dive research into generative AI, including a new Section 7 on the expected change in capital-expenditure allocations in the wake of the DeepSeek disruption. Beginning in Section 3 and throughout this report, our interactive market-sizing models are used to forecast growth potential for total generative-AI spending and segments including, hardware, software, training, inference, advertising and services. These calculators are available on the Terminal at **BI INET <GO>** and **BI GENAG <GO>.**

BI on AI

**March 24, 2025**

# Contents


| | | |
|---|---|---|
| **Section 1.** | **Executive Summary** | **2** |
| **Section 2.** | **Catalysts to Watch** | **3** |
| **Section 3.** | **AI Overview** | **4** |
| **Section 4.** | **Market Disruption** | **16** |
| **Section 5.** | **Segment Analysis** | **21** |
| **Section 6.** | **Expanding Uses** | **39** |
| **Section 7.** | **Hyperscaler Capex** | **53** |
| **Section 8.** | **Processing, Memory Chip Demand** | **65** |
| **Section 9.** | **Regulatory Landscape** | **72** |
| **Section 10.** | **ESG Outlook** | **80** |
| **Section 11.** | **Performance and Valuation** | **82** |
| **Section 12.** | **Company Impacts** | **85** |
| **Section 13.** | **Methodology** | **89** |
| **Section 14.** | **Glossary of Terms** | **93** |
| **Bloomberg Intelligence Research Coverage** | | **95** |
| **Copyright and Disclaimer** | | **98** |
| **About Bloomberg Intelligence** | | **100** |



| **Lead Analysts** | | |
|---|---|---|
| Mandeep Singh | Software, Internet, Hardware Americas | msingh15@bloomberg.net |
| Anurag Rana | Software, IT Services, Americas | arana4@bloomberg.net |
| **Contributing Analysts** | | |
| Masahiro Wakasugi | Japan Technology, Global Semis | mwakasugi4@bloomberg.net |
| Poonam Goyal | E-commerce & Athleisure | pgoyal4@bloomberg.net |
| Steven Tseng | Technology, Hardware, APAC | htseng18@bloomberg.net |
| Sunil Rajgopal | Software, Americas | srajgopal4@bloomberg.net |
| Charles Shum | Semiconductors, Electronic, APAC | cshum2@bloomberg.net |
| Woo Jin Ho | Hardware, Networking, Americas | who88@bloomberg.net |
| Kunjan Sobhani | Semiconductors, Americas | ksobhani@bloomberg.net |
| Tamlin Bason | Tech Policy, Litigation, EMEA | tbason3@bloomberg.net |
| Breanne Dougherty | Equity Strategy - Thematics | bdougherty25@bloomberg.net |
| **Editorial & Visuals** | | |
| Rik Stevens, Justin DeVoursney, Philippe Tardieu | | |


**More detailed analysis and interactive graphics are available on the Bloomberg Terminal**



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Section 1. Executive Summary

**$1.8 Trillion**

AI-driven sales by 2032 from about $93 billion in 2023

**30%**

Projected CAGR through 2032

**$735 Billion**

AI inferencing market

## Racking Up 14-16% of Technology Spending

Generative artificial intelligence is poised to produce $1.8 trillion in revenue, or 14-16% of all technology spending, through 2032 in hardware, software and services and more as businesses supercharge their products and more consumers embrace the shift. The interfaces and tools leveraging gen AI are early, but we believe some common themes include generating summaries, personalized recommendations, contract reviews, image and video content using conversational user interfaces, and built-in language translation. Training AI through machine learning and neural network algorithms using massive datasets (the large language model, or LLM) will be a huge market, reaching $580 billion in sales by 2032 and boosting demand for accelerators on servers and storage units at data centers. Companies will use the public cloud to deploy generative AI, benefiting hyperscalers like Meta, Microsoft, Amazon and Alphabet, with projected compound annual sales growth of 36%, to $236 billion, driven primarily by inferencing workloads.

### Key Research Topics

- **Inference to Surpass Training Earlier:** Compute consumption for inferencing is likely to overtake training workloads at least three years sooner than our initial projections. The biggest driver for the change in our assumptions: The pressure on hyperscalers to generate return on assets for their capital spending and the pivot toward focusing on cost efficiencies for inferencing after DeepSeek's release.
- **Narrowing the Gap Between LLMs:** A convergence in model functionality between OpenAI GPT, Google Gemini, Meta Llama, Anthropic Claude, DeepSeek and Mistral has increased the likelihood that increased commoditization in LLMs with the lowest-cost provider will be based on open-source model and geographic considerations around the storage of data
- **Inference on Applications and Devices:** Smartphone makers such as Apple and auto OEMs like Tesla could reap benefits from the demand for inference-based conversational AI products and vision AI offerings tied to generative AI. Quantization is increasingly being used to shrink the size of trained models for inference on edge devices.
- **High-Bandwidth Memory (HBM) Chips:** As AI models become more complex and training more demanding, HBM chips, which generate high operating profit margins, should be more widely adopted. The speed of chip-performance improvement required for AI is faster than the evolution of miniaturization and advanced packaging, aiding the role of chip testers.

### Performance and Valuation

Year-to-date, our BAIAET Index of AI-themed stocks is down 6%, trailing most BI themes and the B500 Index. The index includes 119 companies, and performance in the first half of last year was driven by Nvidia, other hyperscalers and companies like Super Micro Computer and AMD.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Section 2. Catalysts to Watch

## Model Scaling Driving Adoption

Spending on generative AI has quickly become non-discretionary for enterprises, and we expect sharp growth to be driven by steady hardware investments, uptake of chatbots and attached subscriptions for copilot-type offerings. Retrieval-augmented generation (RAG) allows enterprises to leverage proprietary data lakes to enhance accuracy and reduce hallucinations within LLM responses. Already, companies such as Nvidia have seen huge moves in growth forecasts as a result of the push into AI, while others, like Microsoft (Azure consumption and copilots) anticipate robust gains. On the consumer side, Google Overview was rolled out at scale while Meta has 700 million MAUs for its Meta AI offering.

### Critical Milestones:

- **2024:** GPU, accelerator-chip availability improves for training LLMs. **Surpassed.**
- **2024**: New versions of foundational LLMs based on chain-of-thought reasoning and reinforcement learning. **In progress.**
- **2024**: Strong attach rates for copilots launched by software companies. **Not met.**
- **2024:** AI agents disrupt customer service and help save operations costs. **In progress.**
- **2024:** RAG makes enterprises comfortable with deploying open-source foundational LLMs on-premise. **In progress.**
- **2024:** New content-generating tools, ad-targeting improvement from large internet companies. **Surpassed.**
- **2025:** Multimodal LLMs become a feature across internet apps, enterprise software.
- **2025:** EU on course to adopt first comprehensive regulations through the AI Act.
- **2026:** On-device AI gains momentum with new features for mobile, PCs, AR/VR.
- **2026:** HBM4 could be widely adopted for next generation AI computing.
- **2023-27:** TSMC's generative AI segment reaches compound annual growth of 50%.
- **2027:** AI networking could expand by 5x, driven by specific accelerator requirements.
- **2027-2028:** HBM4E likely to be introduced.
- **2030:** AI agents for contract reviews and customer service hits $115 billion (10% of the total) from $1.4 billion in 2024.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Section 3. AI Overview

## Addressable Markets Appear Ready to Expand

Training foundational LLMs is still the biggest driver of incremental revenue for generative AI, yet traction with GitHub Copilot and growing interest in new apps like Perplexity for consumer search and Sora for prompt-based video generation may continue to expand the addressable market. Generative AI is poised to be a $1.8 trillion market by 2032 (see Figs. 1 and 2), as it boosts sales for tech's hardware, software, services, ads and gaming segments, growing at a compound annual rate of about 30%, based on BI's interactive market-sizing model. As the revolutionary technology changes how businesses operate and enhance their products and services, generative AI could expand to 14-16% of total information-technology spending in such segments from less than 2% today.

**Figure 1: Generative AI Revenue Potential**

($ million, unless otherwise specified)

| Generative AI Revenue Projections | 2023 | 2024 | 2028E | 2032E | 2024-32E CAGR |
|---|---|---|---|---|---|
| **Infrastructure (Training)** | **$75,519** | **$150,179** | **$325,585** | **$579,526** | **18%** |
| AI Servers | $56,642 | $119,325 | $190,373 | $317,419 | 13% |
| AI Storage | $10,893 | $14,916 | $34,267 | $59,055 | 19% |
| Training Compute (Cloud Workloads) | $3,050 | $4,972 | $22,845 | $44,291 | 31% |
| Networking | $3,268 | $5,966 | $19,037 | $36,909 | 26% |
| LLM Licensing Revenue | $1,667 | $5,000 | $59,063 | $121,852 | 49% |
| **Devices and Applications (Inference)** | **$10,459** | **$38,655** | **$285,075** | **$734,732** | **44%** |
| Inference/Fine-Tuning (Cloud Workloads) | $2,179 | $14,916 | $76,149 | $191,928 | 38% |
| Drug Discovery Software | $32 | $90 | $13,571 | $41,203 | 115% |
| Computer Vision AI Products | $2,813 | $5,000 | $27,147 | $63,517 | 37% |
| Conversational AI Products | $3,750 | $13,001 | $70,583 | $114,330 | 31% |
| Chatbots/AI Agents | $1,428 | $4,982 | $90,809 | $311,497 | 68% |
| Customer Service / Contract Review AI Agents | $238 | $1,411 | $58,084 | $213,862 | 87% |
| Cybersecurity Copilots | $21 | $58 | $5,351 | $14,133 | 99% |
| Coding and Devops Copilots | $238 | $1,128 | $13,940 | $37,426 | 55% |
| Educational Copilots | $617 | $1,262 | $6,303 | $20,995 | 42% |
| E-Commerce Copilots | $263 | $532 | $3,340 | $11,880 | 47% |
| Social Media Chatbots | $53 | $591 | $3,791 | $13,200 | 47% |
| Educational Content Creation | $123 | $140 | $700 | $2,333 | 42% |
| Data Protection | $134 | $526 | $6,115 | $9,923 | 44% |
| **Generative AI Driven Ad Spending** | **$4,638** | **$10,799** | **$70,179** | **$208,987** | **45%** |
| Search | $2,472 | $5,522 | $27,959 | $70,174 | 37% |
| Videos | $1,667 | $3,770 | $30,962 | $100,783 | 51% |
| Messaging | $500 | $1,508 | $11,259 | $38,031 | 50% |
| **IT Services** | **$169** | **$8,838** | **$43,416** | **$80,443** | **32%** |
| **Business Services** | **$79** | **$207** | **$15,063** | **$31,097** | **87%** |
| **Workload Monitoring Software** | **$1,216** | **$2,719** | **$20,459** | **$81,603** | **53%** |
| **Entertainment/Media** | **$837** | **$3,353** | **$31,344** | **$86,883** | **50%** |
| Gaming | $522 | $2,644 | $26,109 | $71,042 | 51% |
| Virtual Goods | $130 | $1,058 | $9,791 | $26,641 | 50% |
| Game Design | $391 | $1,586 | $16,318 | $44,401 | 52% |
| Image/Video Generation Tools | $263 | $591 | $4,513 | $13,200 | 47% |
| Film Production/Music Generation Tools | $53 | $118 | $722 | $2,640 | 47% |
| **Total** | **$92,918** | **$214,752** | **$791,121** | **$1,803,271** | **30%** |

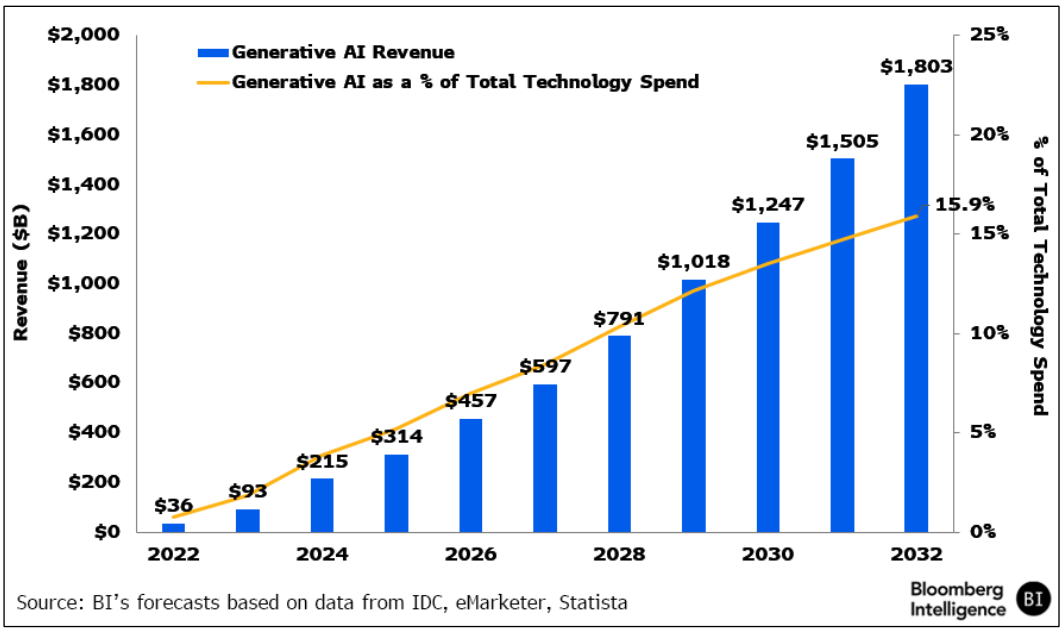
Source: BI's forecasts based on data from IDC, eMarketer, Statista

Bloomberg Intelligence BI



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 2: Generative AI Spending**




BI

**$1.8 trillion in spending seen by 2032**

Semiconductors, hardware, cloud software, IT services and advertising companies may be the vanguard of these changes. Yet we also could see new products and services, a displacement of incumbents and the emergence of new categories.

## 3.1 Inference to Surpass Training Spend with Reasoning

The success of reasoning and chain-of-thought LLMs has paved the way for the inference market to exceed that for training by 2029, at least three years sooner than our earlier view. Increased deployment of AI agents could accelerate the demand for inference AI GPUs and other accelerator chips as tokens processed by reasoning models will be an order of magnitude higher than current pre-trained models. Agentic AI, by nature of its autonomous task performance, will require visual perception models in addition to text-based models. Training could be about 30-35% of added revenue and nearly a $580 billion market by 2032, encompassing servers, storage, networking, service offerings on the cloud, as well as LLM licensing. By 2032, 26% of the $735 billion inference market will be related to cloud workloads, growing faster than the overall infrastructure market at a 38% compound annual rate. The high growth expectations have so far been supported by the rapid change in data-center compute and storage to handle AI workloads across both consumer and enterprise apps. Nvidia has been the most significant beneficiary of the trend, while hyperscalers have all boosted capex to ensure GPU availability for both internal consumption and their cloud businesses.

As for hardware, computer vision and conversational AI products may emerge as new categories with the availability of LLMs for domain-specific predictions and applications. Such new items can help accelerate growth in the $1 trillion devices market, which has blossomed with smart speakers and wearables.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 3: Training vs. Inference Forecasts**

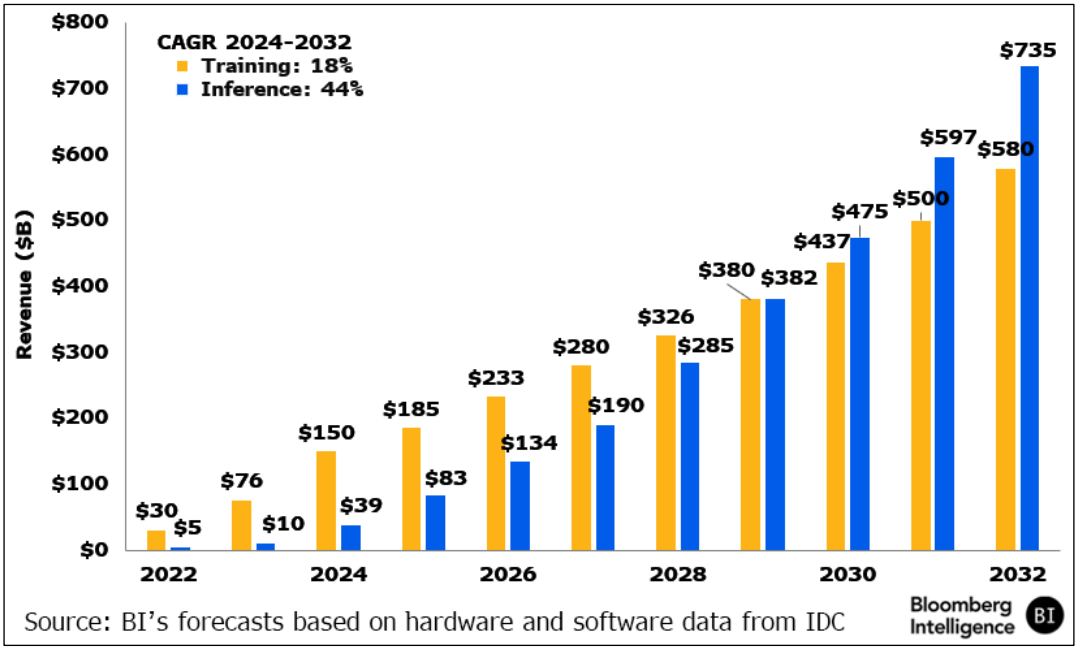
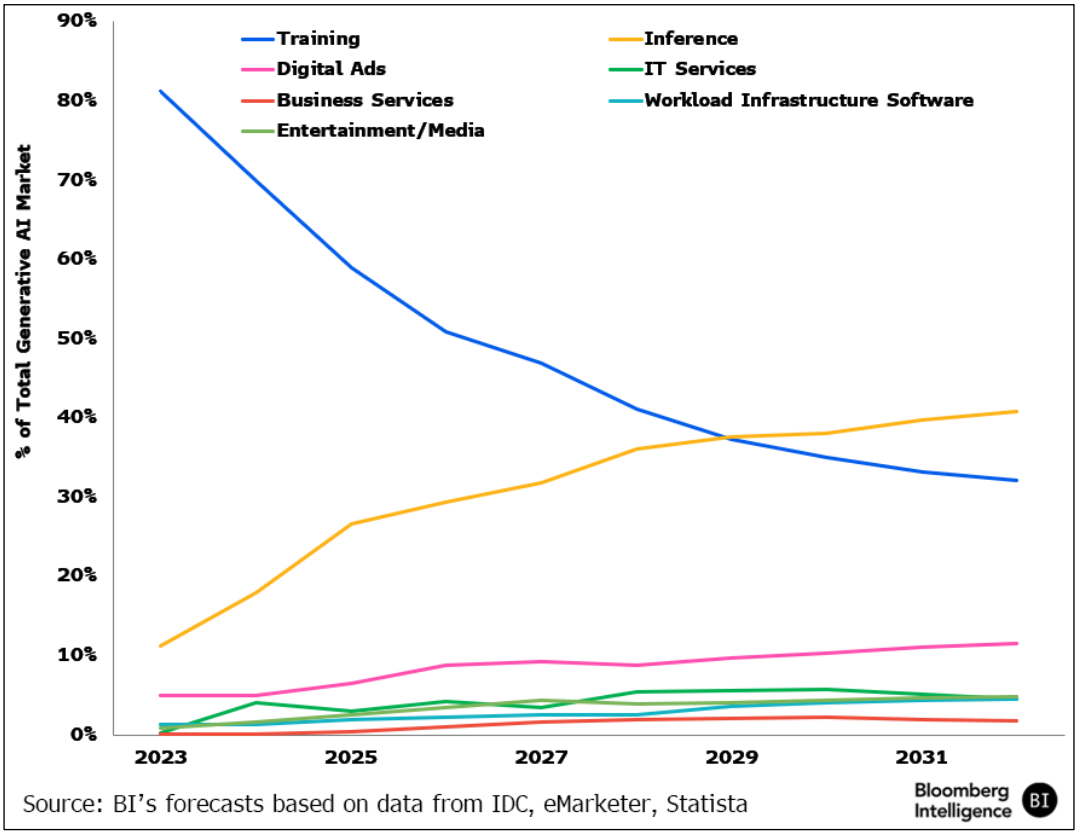



**Figure 4: Generative AI Market Share**

BI

**Inference outlays exceed training three years earlier than initially projected**






No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

Within hardware, infrastructure expenditures (for training) will likely be about triple the size of those for devices (inference) as companies spend on servers and storage to manage the intensive workloads required. Nearly 75% of chief information officers in the US surveyed by Bloomberg Intelligence indicated plans to increase their IT-infrastructure budgets in 2024, with about 38% expecting to boost spending by over 11%. Nvidia and Dell may remain among the top server providers for generative AI workloads, with about 51% of CIOs selecting them. Generative AI infrastructure as a service (IaaS) will be key to training LLMs and could add $236 billion in sales over a decade. Within generative AI IaaS, the market for compute resources for training will likely be the largest at around $44 billion, while inferencing workloads run on the cloud could result in a $191 billion market through 2032. Networking can add about $37 billion The market for computer-vision AI products is set to grow to $64 billion, while sales of conversational AI products could hit $114 billion. We expect that AI may add $580 billion to the total AI training market by 2032 from roughly $75 billion in 2023.

In software, generative AI products may add about $446 billion in spending by 2032, growing at a compound annual rate of 64%, withcybersecurity, drug discovery, AI assistants and coding workflow among the top beneficiaries. Many software peers will likely introduce their own AI copilots to enhance the user experience, with specialized assistant software poised to log $95 billion in sales by the end of the decade. Spending on educational software could be strong in an effort to improve existing learning tools and build new ones. We also expect generative AI to expedite development of gaming and creative software, reducing barriers to entry and creating opportunities for disruption.

On the internet side, generative AI can improve ad targeting and spur the creation of new formats to drive user engagement and increase conversion of ad views to sales. Large companies like Meta and Alphabet rely less on open-internet corpuses than other companies developing foundational LLMs, given their abundance of first-party data to deploy, along with robust capacity to spend that can help train models to improve ad targeting and efficiency. Such enhancements may drive an additional $210 billion for the digital-ad sector by 2032.

In IT and business services, we estimate that generative AI products and tools can add about $111 billion in sales as companies look for new products to drive top-line growth and trim unnecessary costs.

## 3.2 Cloud to Overtake Server Development

Though servers and storage could be the most prominent segments for generative AI services in the near term, many enterprises doubtlessly will leverage public cloud deployment eventually. We believe hyperscalers will develop in-house foundational LLMs, which will work best on their own cloud infrastructures. Meta, Microsoft, Alphabet, Nvidia, Amazon and other such suppliers may be among the main facilitators for training LLMs. These companies have access to the capital needed to set up the training infrastructure while keeping usage high for their servers to sustain healthy profit margins.

In time, generative AI as a service should be a much bigger market than servers and storage, as seen in Fig. 1, logging 13% growth compounded annually through 2032 as gains for stand-alone servers and storage taper off. The trend favors expansion for hyperscale cloud suppliers over


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

smaller infrastructure-software peers, mirroring the evolution of the software-, platform- and infrastructure-as-a-service portions of the roughly $663 billion public cloud market.

**Figure 5: Generative AI Market Overview**




*Source: Bloomberg Intelligence*

BI

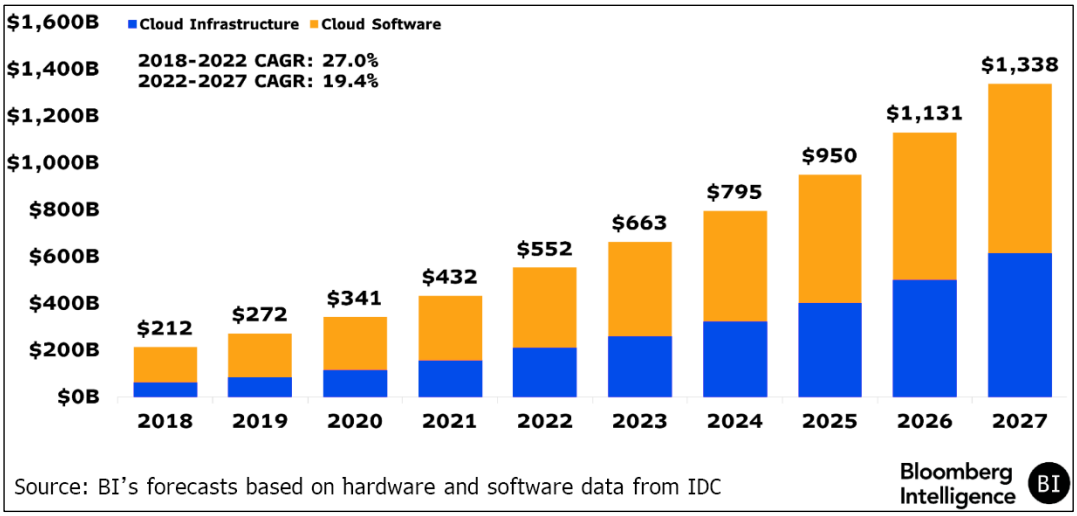
**From drug research to gaming, a wide array of uses for AI**

The forecasts in the accompanying graphic are conservative. Though it's highly likely that the enterprise demand shift toward the cloud will gather speed in the coming years, that isn't included in our assumptions.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

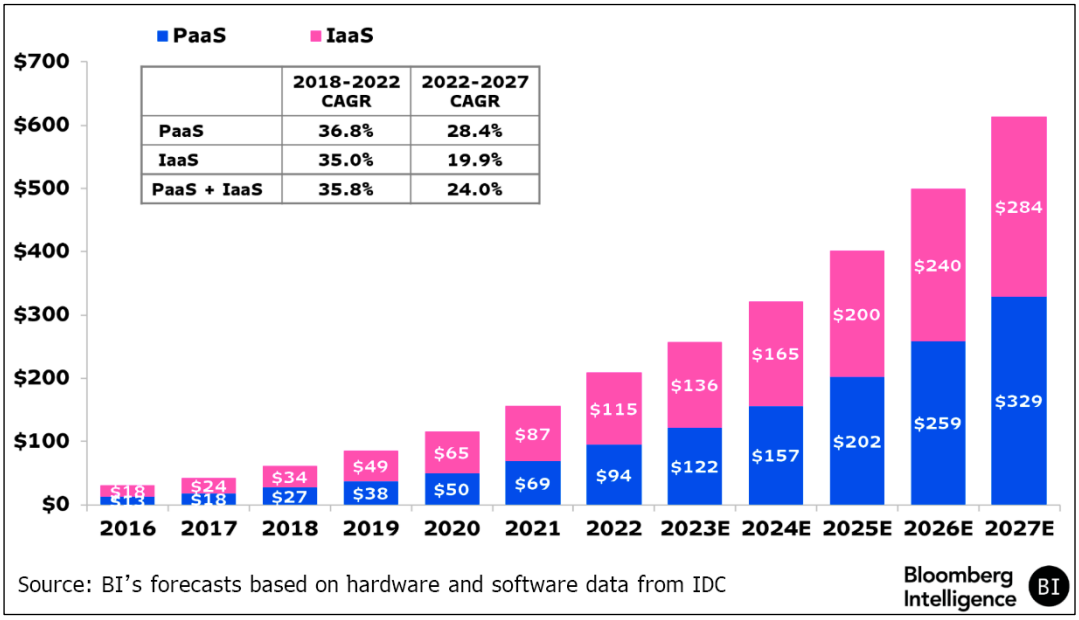
**Figure 6: Total Public Cloud Spending Forecast ($ Billion)**




BI

**Public cloud spending heads for $1.3 trillion, Iaas $284 billion, PaaS $329 billion**

**Figure 7: IaaS, PaaS Revenue Forecast ($ Billion)**




## 3.3 Near-Term Server Demand Should Stay Healthy

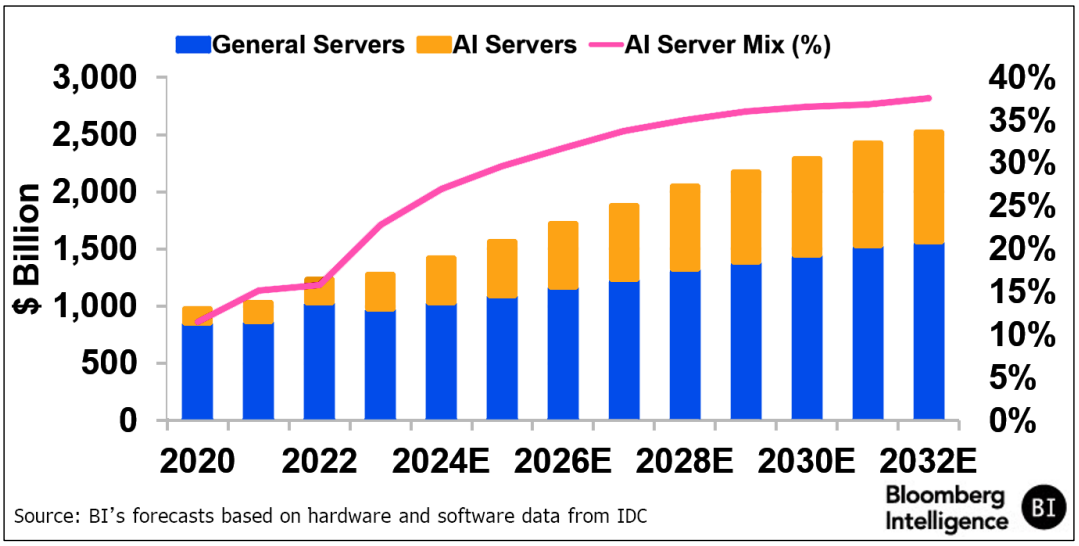
The eventual shift to cloud deployment notwithstanding, the explosive demand for generative AI – as evidenced by the insatiable demand for Nvidia GPUs – should fuel significant growth in the infrastructure hardware market, particularly servers (Fig. 8), that provides the necessary computing power. The global market for AI servers is set to roughly double to $39.4 billion in 2024 from 2022, according to our estimates, notching 41% average annual growth. AI is poised to contribute more than 20% of global server revenue starting this year, from 15% in 2021. Despite



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

economic headwinds in 2023, spending on AI servers could remain robust, thanks to the ChatGPT-fueled arms race in generative AI.

**Figure 8: Worldwide AI Server Market Forecast**




BI

**A growing share of the server market**

The lion's share of server demand might go to original design manufacturers (ODM) building customized models for major cloud service providers like Microsoft and Google that are providing significant backing and development for AI applications. Their public cloud infrastructures also offer the necessary scalability for AI development in terms of computing and storage capability. Microsoft is a key investor in OpenAI, the owner of ChatGPT, and Microsoft Azure is the exclusive cloud platform for ChatGPT. After posting triple-digit growth in AI server revenue in 2024, major ODMs in Taiwan – Hon Hai, Quanta, Wistron/Wiwynn and Inventec – expect this momentum to extend into 2025 given the steady launches and ramp-up in shipments of new AI server projects.

The complex designs of AI servers may help boost profit for related suppliers. While typical servers mainly use Intel and Advanced Micro Devices' x86 central processing units, AI servers specialize in parallel processing, which requires accelerator chips such as graphic processor units (GPUs) and advanced RISC machine-based (ARM-based) chips — or proprietary application-specific integrated circuits, such as Google's tensor processor unit (TPU). Whereas a traditional server has many as two CPUs, AI servers can have up to eight GPUs, each valued at $30,000-$40,000, leading to an eight to 10-fold pricing premium for AI servers over traditional servers.

The need to optimize system performance, power efficiency and manage cooling poses a challenge to server designs as each processor has different instruction sets and data transmission cycles. As a result, ODMs with advanced design expertise would have a competitive edge over rivals that don't and might be able to charge more, expanding profitability. It's worth noting that AI server gross margins may be lower than regular servers, as higher gross profit per unit is proportionally smaller relative to significantly higher average selling prices. That said, for server ODMs, operating margins can still be better with sufficient shipment scale.



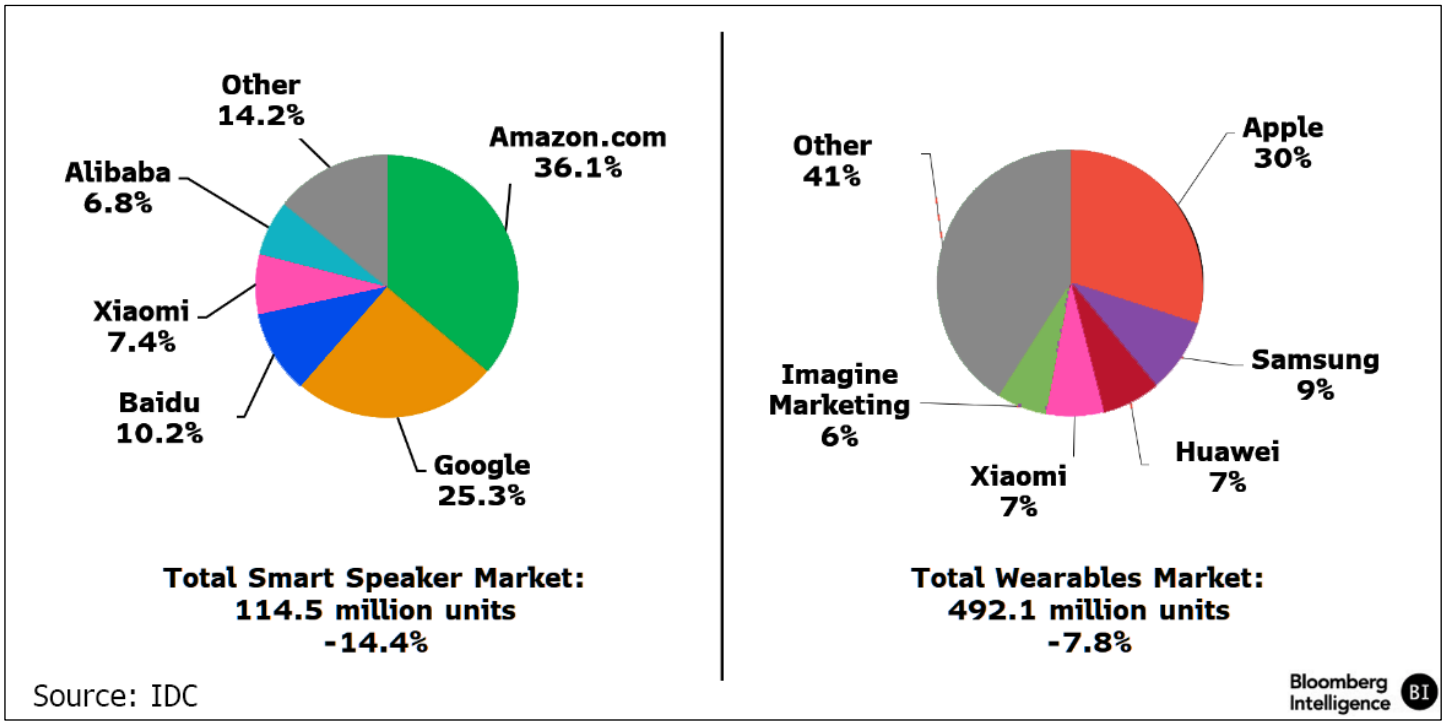
No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

The shift from using general-purpose CPUs to custom accelerators for large dataset workloads is key to why training is poised to be a bigger market (accounting for 70% in 2030 and 63% in 2035) than devices (20-30%) in generative AI. The use of semiconductor accelerators will likely increase as more companies train their own LLMs, similar to OpenAI's GPT, Meta's Llama and Alphabet's Gemini.

## 3.4 Faster Hardware Refreshes; Networking a Key

The need for inferencing on so-called edge devices (hardware that controls data flow across the boundary between two networks) could speed refreshes of personal computers and smartphones – which currently aren't well-suited for the heavy processing, memory and storage requirements for AI's LLMs – while spawning new categories beyond wearables and smart speakers. Demand for inference is expected to ramp up as more applications are developed on top of foundational models such as OpenAI's GPT, Google's Gemini and Meta's Llama. We expect more compact versions of these models, such as the Gemini Nano, to be released in the next 1-2 years, allowing edge devices to perform AI workloads natively. These models will likely require less computational power and be less expensive to train.

**Figure 9: Smart Speakers, Wearables Market 2022**

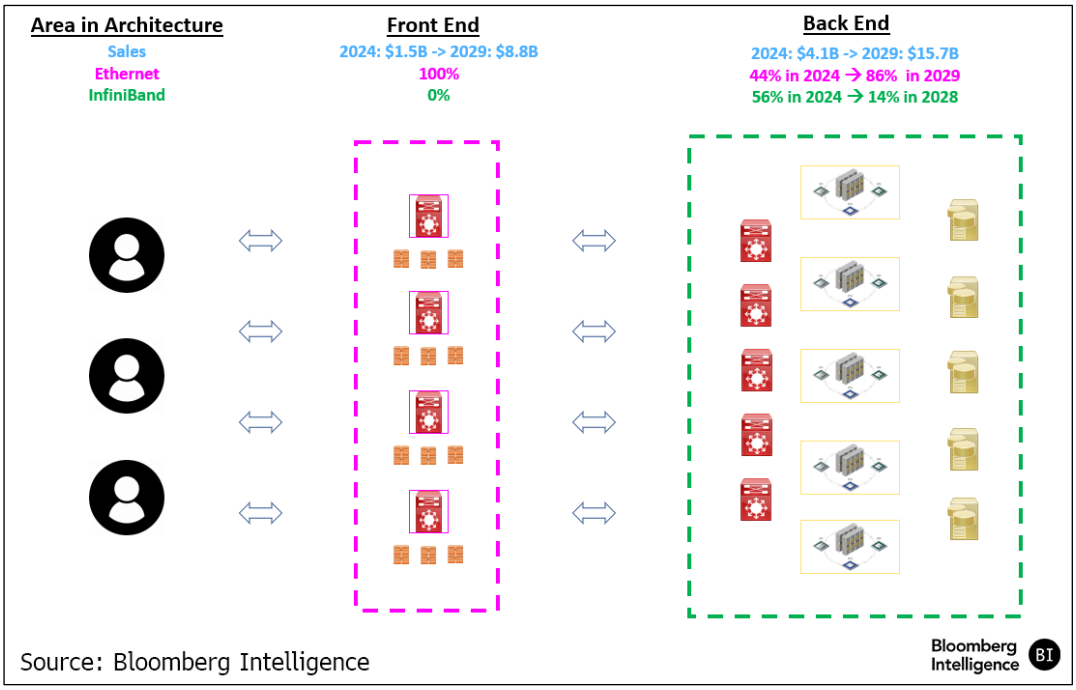



Networking has emerged as a strategic area of AI infrastructure, along with servers and graphics processing units. It's typically been a bottleneck for hyperscale cloud infrastructure, which companies have aimed to resolve with higher capacity gear. Public cloud generative AI workloads are expected to grow faster than general cloud due to the computing intensity needed to ingest


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

a burgeoning amount of structured and unstructured data for LLMs, supported by clusters of GPU-powered servers that can number in the tens of thousands. This complexity and density of AI architecture requires a separate "back end" AI network that supports high capacities – 800 gigabits or greater – while delivering low-latency, error-free data traffic separate from the general purpose, user-facing "front-end" networks. Given the rapid rise of AI architectures, backend networks are expected to account for roughly 25-30% of total Ethernet data center ports by 2029.

**Figure 10: AI Network Architecture Overview**




BI

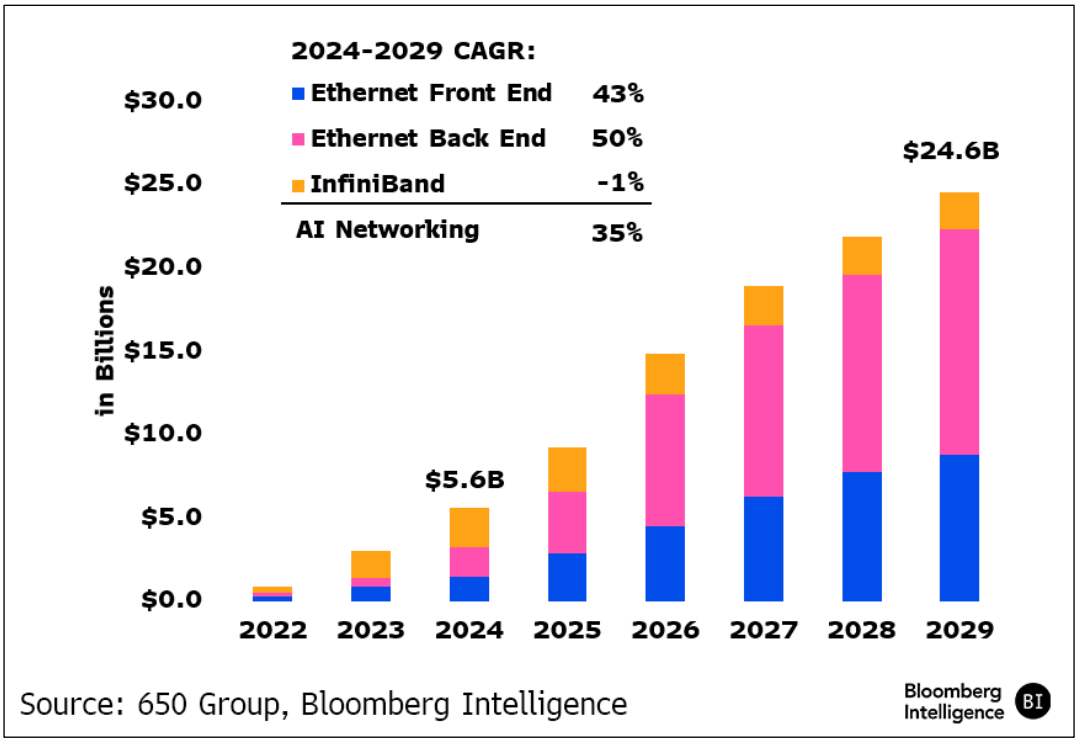
**Architecture nears $12.5 billion in 2028**

AI backend networks are supported by two distinct protocols, InfiniBand and Ethernet. InfiniBand emerged as the preferred connectivity due to its lineage, high-performance computing and supercomputing environments, which is due to the ability of the technology to reliably transport data at high speeds with little data loss. This allowed the technology to account for 73% of AI back-end networking sales in 2023.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 11: Cloud AI Networking Sales Forecasts**




BI **AI networking sales could grow at 31% a year through 2028**

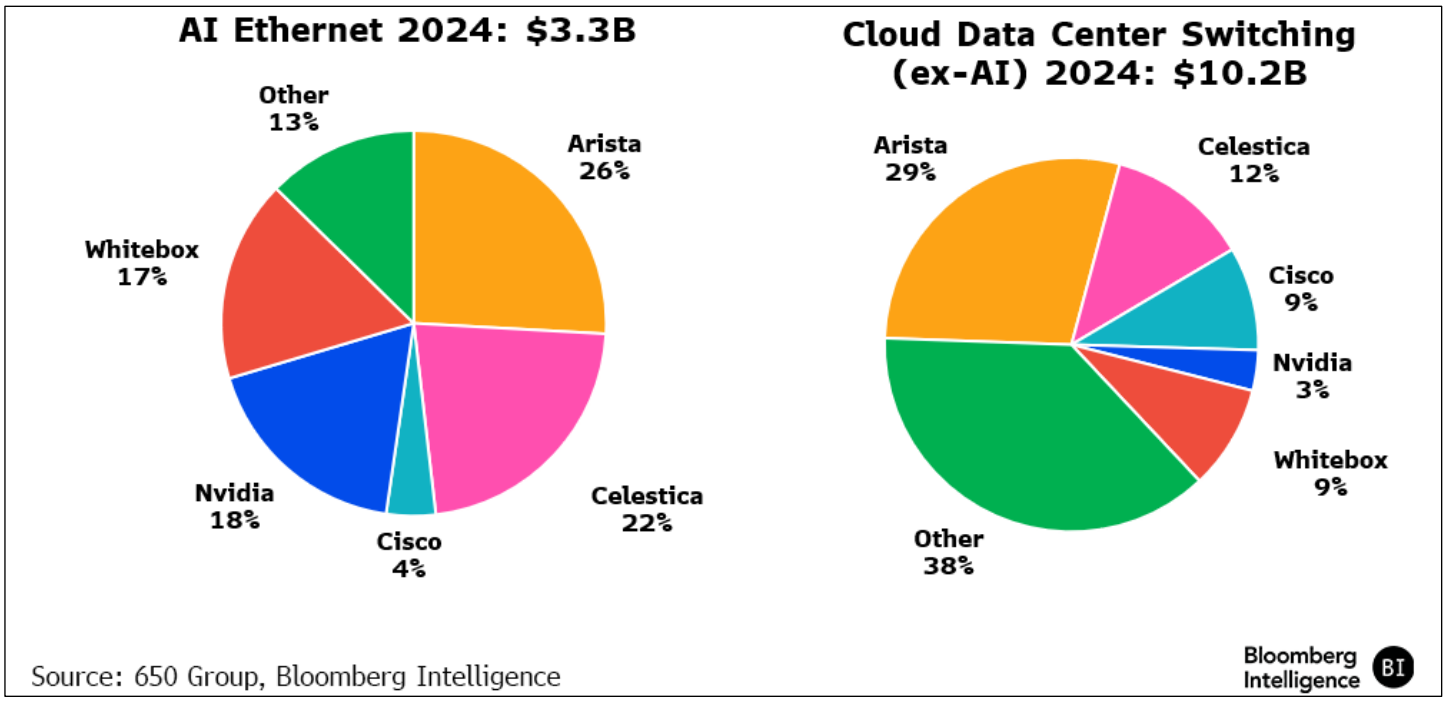
However, cloud customers are rapidly building out their AI infrastructures on Ethernet networks, helped by protocol and chip enhancements and a growing ecosystem of products led by the Ultra Ethernet Consortium. As a result, back-end AI-related Ethernet sales are expected to grow at a 50% compound annual rate for 2024 to 2029 to $13.5 billion. Hyperscale cloud interest in adopting Ethernet may be high in part because of familiarity with the technology but also because it allows each cloud to create differentiated AI architectures and helps avoid lock-in with the Nvidia ecosystem. Arista's strength in high-speed networking gear positions it well to be a leading beneficiary of the shift toward Ethernet by cloud providers.

InfiniBand will still play a meaningful, but diminishing role in AI networking, as tier-2 and emerging neo-clouds may leverage Nvidia for a turnkey, bundled AI solution. Nvidia meanwhile is well situated to shift customers to Ethernet for AI networks, thanks to the Spectrum switching gear it gained through its Mellanox acquisition, but this is less likely for hyperscale cloud customers.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 12: AI Ethernet, Total Cloud Switching Market Share**




## 3.5 Digital Transformation Initiatives Spill Over

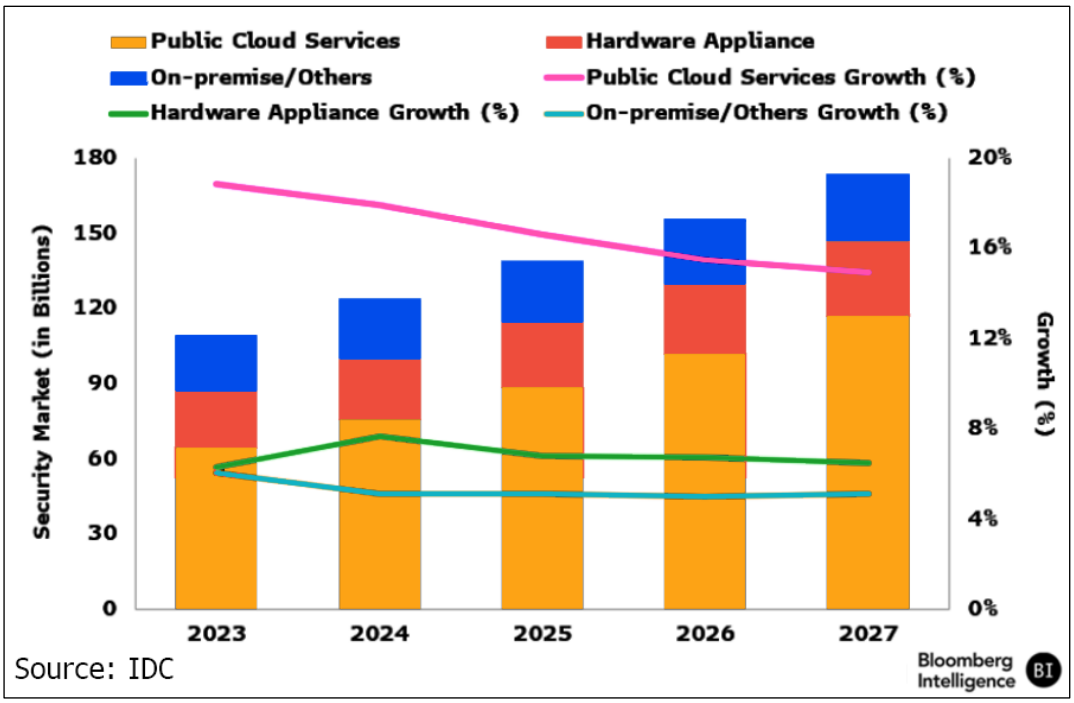
With the rapid development of cloud-based AI technologies like ChatGPT, the significance of edge AI development is growing and represents a significant step in advancing the AI ecosystem. Edge AI is particularly valuable for real-time decision making and cost savings, which are crucial in areas such as health care, manufacturing and transportation and could lead to a larger user base than cloud-based AI. Our scenario analysis finds that the edge AI semiconductor market could be as much as 3.37x the size of the cloud-based AI market by the end of 2032. Adoption of edge AI can drive significant growth in uptake by the consumer (projected to lead other segments with a 10-year CAGR of 39%), industrial and automation sectors over the next 10 years.

Beyond generative AI, advancements in machine learning and other aspects of artificial intelligence appear likely as well. Oracle has been pushing its autonomous databases for the last few years, which could gain from increased budget allocations to AI. We expect increased availability of such features from other software providers, too, where product patching, security updates and tasks usually assigned to a database administrator are automated using machine learning. It could play a much bigger role in cybersecurity in the coming years, as well, especially in event management, in analyzing irregular patterns inside organizations and in the form of copilots to augment security professionals and boost automation.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Public cloud services capture largest share of security market**

**Figure 13: Market Size by Deployment Type**






No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Section 4. Market Disruption

## A Shift Coming in Hardware, Ads, Gaming

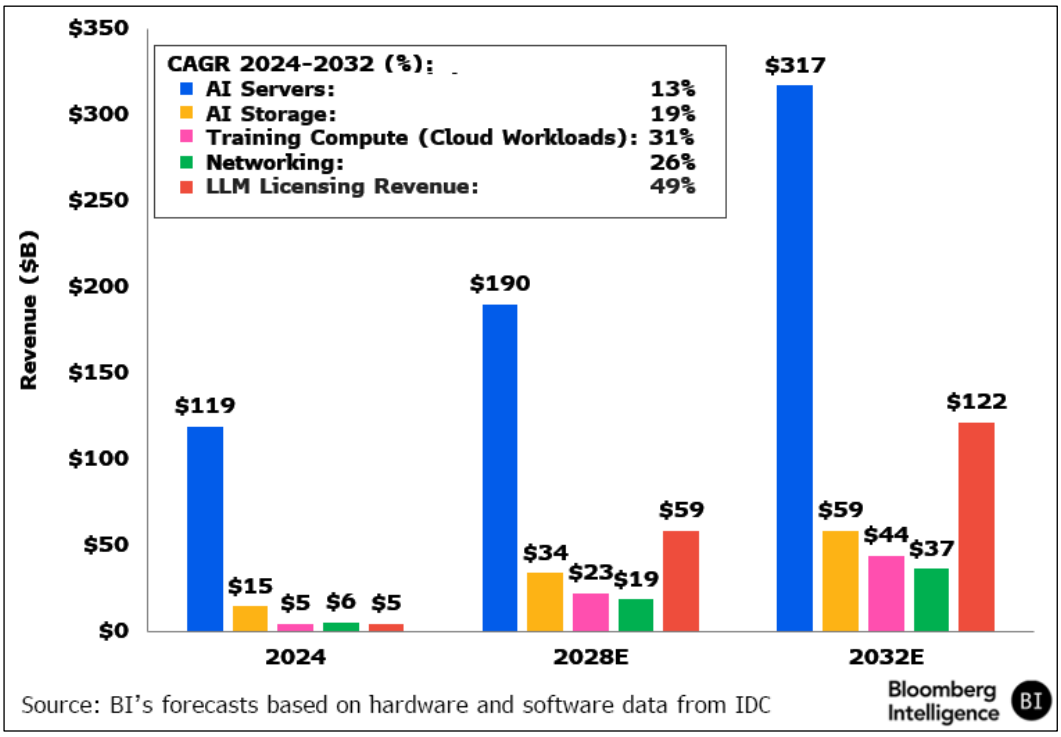
Generative AI offers opportunities for disruption, especially within hardware, digital advertising and gaming. The computational intensity of training large language models may spark a market-share shift toward advanced RISC machines, potentially making this category the fastest growing within hardware. Alphabet, Meta and other digital-ad giants can improve targeting and brand conversions by implementing machine-learning models based on their vast library of first-party data. Sony, Google, Unity and others in the gaming segment may leverage AI to facilitate development and make the user experience more engaging.

## 4.1 LLM Training Favors Move to Accelerators

Training LLMs remains a significant driver for the generative-AI market as demand for AI servers and storage units at data centers grows (see Fig. 14). Nvidia may surpass $200 billion in data-center chip sales in 2025, driven largely by hyperscale customers like Meta, Microsoft and Alphabet. Those companies have said they aim to boost capital spending in the next 1-2 years to meet demand. Companies are leveraging public cloud for deploying LLMs, and we calculate that training compute on cloud could reach $44 billion in spending by 2032. Alphabet and Meta, along with foundational model companies, may see incremental revenue from licensing their LLMs to smaller enterprises that lack the resources to build their own.

BI **Licensing LLMs grows nearly 50% a year**

**Figure 14: Generative AI Training/Infrastructure Spending Forecast**






No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 15: Generative AI Software Spending Forecast**

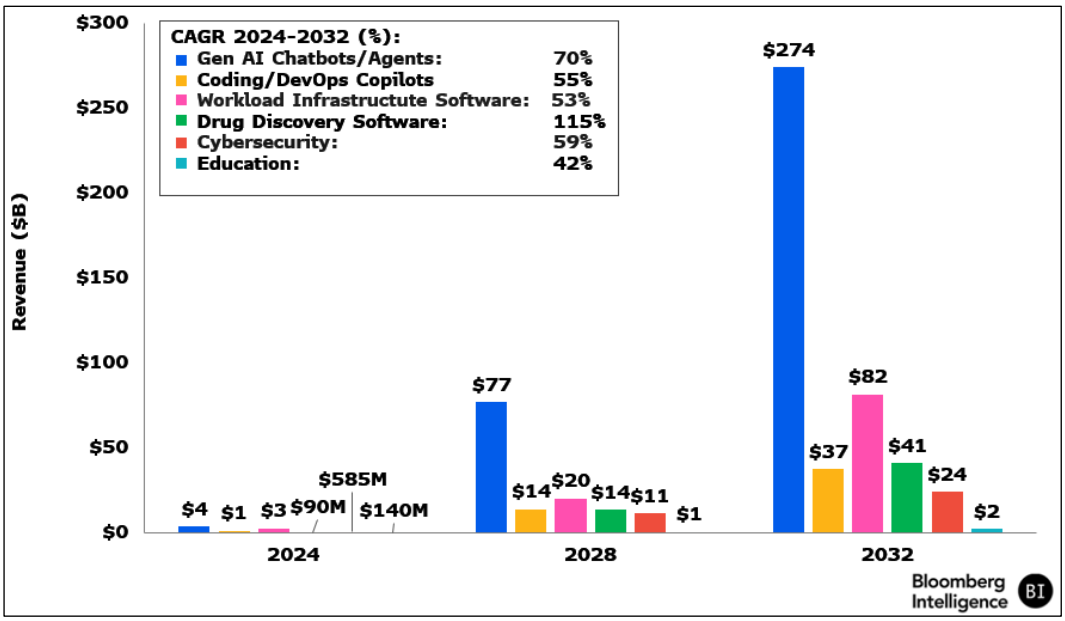
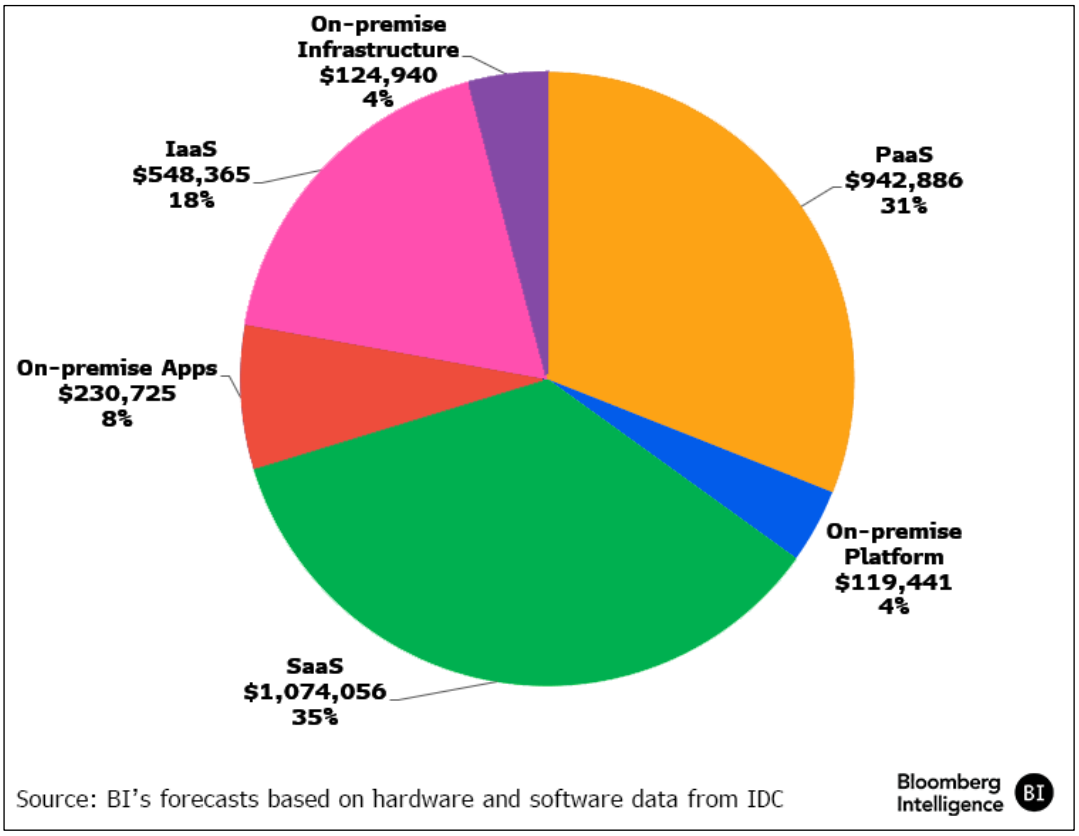



**Figure 16: Software Spending Forecast Breakdown, 2032**

BI

**Software spending growth at 8.3% by 2032; PaaS, SaaS are market leaders**




LLMs have extensive computing and storage needs, central reasons that we expect the first phase of experimentation to be performed with hyperscale cloud providers like Google, Microsoft and



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

Amazon Web Services. Even at maturity, such companies are likely to gain the most market share, given the scale and cost needed to develop an infrastructure in-house.

OpenAI's lead in training transformer models and its early partnership with Microsoft has aided ChatGPT's adoption over text-based LLMs from Meta, Amazon, Alphabet and Anthropic. With most hyperscalers investing in developing their own foundational LLMs, we believe OpenAI will need to maintain its lead on the algorithm side while ensuring access to training data from the open-internet corpuses from companies such as Wikipedia, Reddit and Stack Overflow. Alphabet's merger of its DeepMind and Google Brain AI units, can spur faster changes that leverage LLMs to maintain user engagement across revenue generators, like its Search, Chrome and Maps applications. Amid the rapid shift to generative AI, large players like Meta, Adobe, Microsoft, Alphabet and Salesforce are better positioned than smaller rivals for two reasons: they have scads of first-party data in hand and ample ability to deploy capital. Each has a leading market share in its category, offering access to large amounts of information to train AI models, driving more accurate and efficient results.

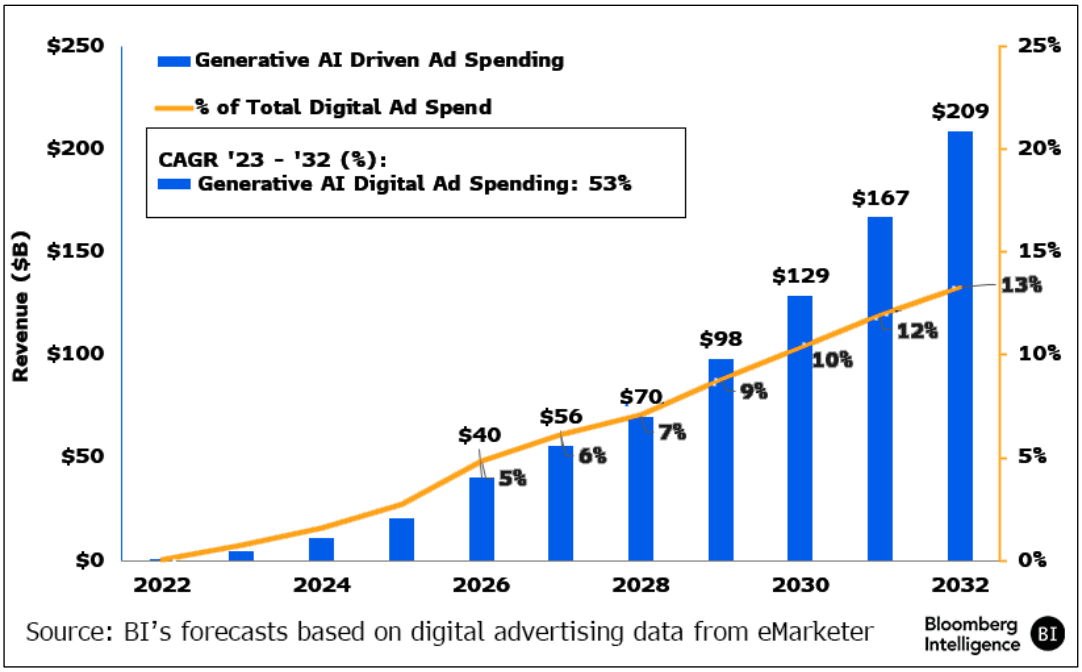
Social-media platforms like Meta should see a lift as AI-generated content rises rapidly, helping to boost engagement and monetization. LLMs and generative AI can accelerate the shift to digital ads from linear TV. We calculate that more time spent online, coupled with ad targeting and personalization could add over $209 billion to the market through 2032 (Fig. 17). Conversion rates for ads on these platforms might be aided by the growing capabilities of LLMs, which should favor companies with strong presences in cloud infrastructure and that have the most first-party data.

BI **Most say they won't pay more than $10 a month for ChatGPT**

Given the high costs of generative AI infrastructure, OpenAI, Anthropic and Google currently have paid subscriptions for their LLM offerings for heavy users. An ad-supported model is unlikely to be profitable for online search and new tools that leverage deep learning and generative AI due to the higher costs of LLM-based queries. A recent Bloomberg Intelligence survey found that only 13% of respondents were willing to pay for a subscription to use a generative AI tool like ChatGPT. Of those, just 1% said they would spend $20 a month for ChatGPT, with the rest amenable to $6-$10. Among all participants, 93% indicated that they wouldn't lay out more than $10 monthly. The results suggest that lower prices could help boost adoption by 10x for a generative AI subscription. As an illustration: though a freemium version of ChatGPT helped propel it to 100 million monthly active users faster than any consumer app at the time, its conversion to paid subscribers remained in the low single digits.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 17: Generative AI Digital Advertising**




BI **Gen-AI ad spending reaches 13% of total digital advertising outlays by 2032**

## 4.2 Sony, Google, Microsoft Parlay New Interfaces for Gaming

Some startups have already showcased interesting techniques to produce synthetic content — computer-made data that mimics real-world information — based on descriptions and the large amounts of training data available on the open internet. Generative AI can speed that creative process for mobile games, social media and virtual- and augmented-reality applications.

AI tools may rapidly increase gaming data available beyond the high-budget, high-profile major publishers, including those made by users. Developers remain key to gaming and the metaverse beyond the foundational models provided by tools such as OpenAI's GPT, Google's Gemini, Meta's Llama and Anthropic's Claude. Apple and Google's Android, as well as gaming ecosystems like Sony's PlayStation could offer software development kits to leverage LLMs to ease creation of new content on their platforms. Generative AI might help creative software tools shift to description- and voice-based content from point-and-click user interfaces. Microsoft's open-source generative AI model Muse, trained on seven years of human game-play data, has the ability to create visuals and simulate sequences of game-play actions, potentially enabling game developers to expedite prototyping process. Individual gamers, too, could prompt the model with images to test out different game-play sequences, with the option to iterate the AI-generated output using their Xbox controllers.

Though Google and Meta have developed LLMs for image generation, they have trailed Stability AI, Midjourney and OpenAI's Dall-E in terms of adoption. OpenAI just released its video-based LLM, Sora, as foundational models become multimodal. Most image-based generative models rely on a diffusion technique and the quality of images rendered depends on training data and

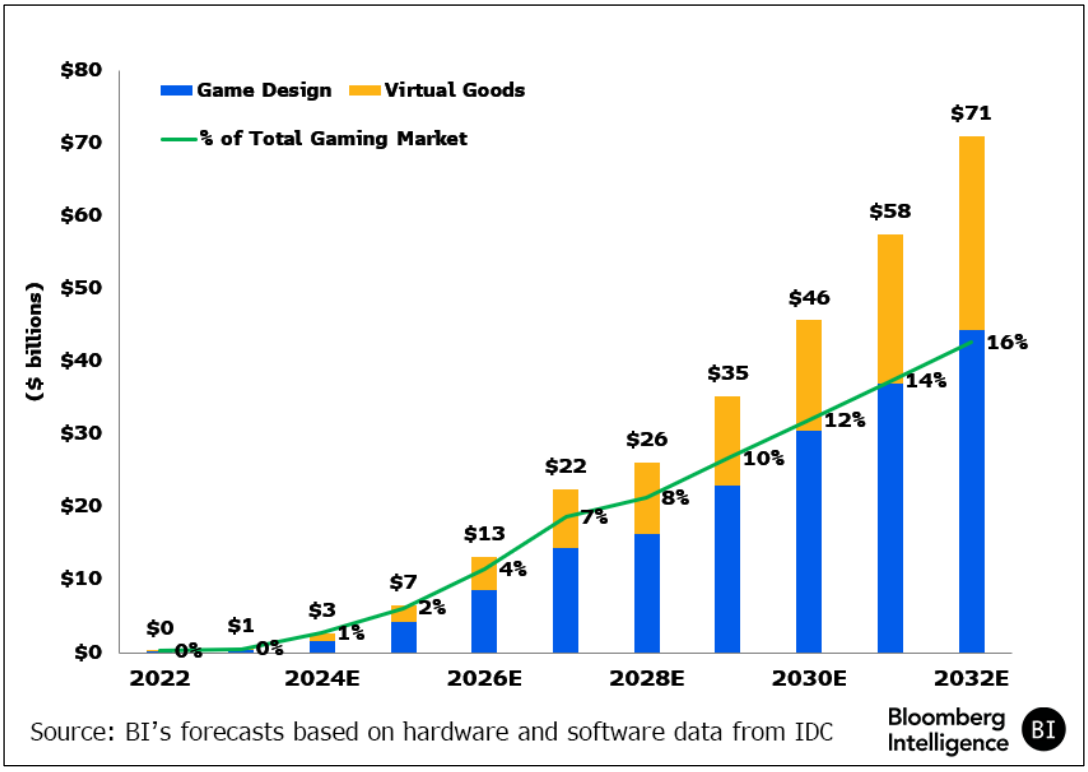


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

the weights assigned to the parameters used. While Adobe has been investing in developing its own generative AI capabilities with the introduction of its Firefly offering, we expect other design- and gaming-software companies to invest in their own generative AI models to leverage proprietary data and distribution.

**Figure 18: Generative AI Gaming**

BI **AI can expand gaming beyond high-profile publishers**






No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Section 5. Segment Analysis

## Popularity Hinges on Adoption in Existing End Markets

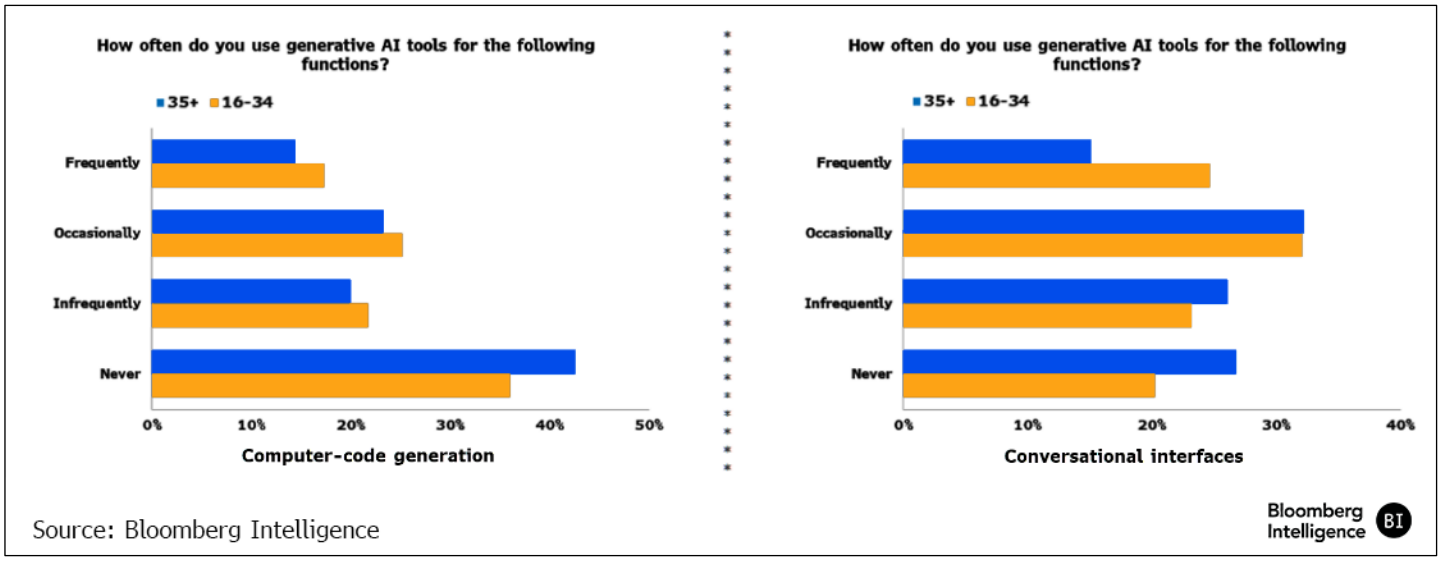
Generative AI is likely to have a far less pronounced effect on the application software industry than infrastructure software in terms of new revenue generated. Within application software, however, we're already starting to see the rise of AI copilots, with companies like Microsoft, Adobe, Snap and others introducing their own versions of this technology in recent months. Bloomberg Intelligence market opportunity analysis shows that a majority of the $309 billion of new software sales tied to generative AI is likely to fall in the infrastructure bucket.

## 5.1 Copilots Lead the Way Into New Endeavors

Education, drug discovery and specialized AI assistants could be the bigger contributors to new revenue streams within application software. Gaming, IT and business services may be smaller categories. The customer-service and business-process outsourcing subsegment of business services might be heavily affected by AI tools and sales could shrink.

Microsoft, Salesforce and Adobe stand out as some of the early adopters of enterprise-AI copilots and agents that focus on sales, customer service and content creation, which Bloomberg Intelligence calculates could grow into a $83 billion market by 2032. Companies with the largest market shares and capital are likely to garner a bigger share of this pie, widening the gap over smaller vendors.

**Figure 19: Coding vs. Conversational Interfaces**






No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

The generative-AI copilot and agent market, consisting of AI chatbots that can assist in various ways, from customer service to customer prospecting and content generation, was at roughly $1.1 billion in 2023, and our analysis of its end markets suggests this could grow 61% a year to $83 billion by 2032 (see Fig. 20). Customer service will likely be the largest use case, expanding about 75% a year to $37 billion, given the more straightforward application of copilots into a primarily chatbot setting.

We view "agents," or autonomous copilots that can now take specific actions, as the next evolution of AI assistants, which could accelerate the market further by removing the need for human intervention for certain tasks, thereby increasing return on investment for customers.

Early adopters in enterprise software of these gen-AI chatbots, such as Microsoft and Salesforce, should stay atop peers in the midterm based on their leading market shares, large customer bases and ability to deploy vast resources to consistently develop new capabilities. Their large customer bases not only provide better upselling opportunities, but also create larger datasets for training and fine-tuning the foundational models that power their chatbots.

BI **Large customer base means more data to train models**

Adobe is uniquely positioned in creative software, given its vast stock image and video library, which it uses to train its content-generation models. We expect this to be one of the hardest competitive moats to overcome for competitors, as this provides Adobe clients with copyright-safe image and video generation.

As AI-copilot adoption rises, we expect companies to price these either on a subscription or consumption-based model, or leverage these tools indirectly to fortify retention and justify price hikes. Applications designed primarily for assisting a user, such as document summarization or email generation, could be more apt for a subscription plan, such as Microsoft 365 Copilot, whereas more compute-intensive uses may lean on consumption pricing. Examples of the latter may include video generation or customer service inquiries.

Generative AI has also led to broad-based subscription pricing increases in cases like Adobe who lifted pricing for all Creative Cloud plans by roughly 9-10%. In other cases, such as ServiceNow, pricing for its generative AI Pro-Plus SKU was based on customer savings due to AI, in which ServiceNow only reclaimed 10% of the perceived customer savings, equating to roughly a 60% price uplift over their Pro SKU.

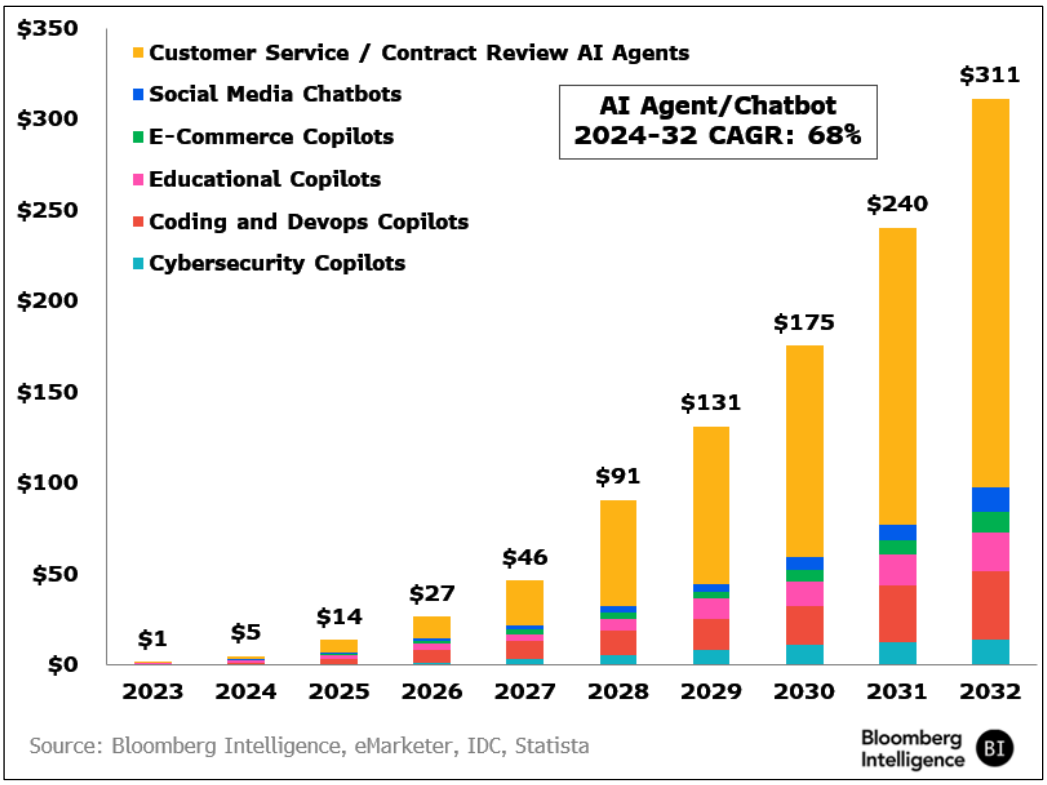
Software providers like Workday have also taken the approach of embedding most gen-AI tools in their products at no extra cost to the customer. We view this as a good way to boost client retention that could lead to future subscription-price increases.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

BI

**AI agent spending from $1.14 billion in 2024 to $214 billion in 2032**

**Figure 20: Generative AI Consumer-Facing Copilot Software Spending Forecast**




## 5.2 Apple Intelligence

Apple's use of its own on-device foundational model that can identify information on the user's screen and take related actions shows the company wants to maintain control of inferencing on its devices. At its WWDC event, Apple showcased App Intents and a semantic index framework for on-device execution, which we believe will be mostly used for tasks like email summaries, image editing and calendar scheduling. Given the widespread use of search, maps and YouTube videos in terms of daily time spent on iOS devices, we believe a large language model integration with Google Gemini would enhance the user experience. Among Android-based smartphone makers, Samsung is utilizing Gemini Nano for on-device AI and Gemini Ultra for larger models on the cloud. The company also showcased its Private Cloud Compute infrastructure for certain types of tasks and an integration with ChatGPT for more complex queries.

Apple showcased AI features such as visual intelligence at its iPhone 16 launch, highlighting AI use cases that are driven by Siri's on-screen awareness, similar to Alphabet's Pixel demo last month. Though Apple is employing small AI models using a semantic index framework vs. Google's Gemini large language model, we believe Apple's iOS 18 software is likely to further narrow the feature gap with its Android rival.

Apple Intelligence will boast a number of new features within iOS 18 as it is released in stages, including text summaries for emails and messages powered by on-device LLMs and an image generator based on prompts, which is already available on Android OS used for the latest Pixel



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

and Samsung devices. The release of Apple's visual intelligence will help the company narrow the gap around on-device search features powered by gen-AI models, as well as potentially increase Siri usage to better compete with Gemini Live Assistant on Android devices.

## 5.3 Enterprise CIO Survey

Though US corporate IT-infrastructure utilization remained unchanged, the use of large language models for copilots and chatbots and investment in equipment like servers were the most significant driver of enterprise spending on IT infrastructure in our CIO survey. Cloud was the top spending priority for 37% of respondents, up from 30% six months ago. This could aid original design manufacturers like Quanta, Hon Hai (Foxconn) and Inventec, which supply customized servers to cloud service providers, particularly Amazon.com and Microsoft.

Some 60% of respondents said they were working on deploying generative AI copilots, almost double the 32% in last year's survey. A lack of proprietary data and customer efforts to hold down costs could limit the adoption of large language models offered by smaller companies and startups, while DeepSeek's open-sourcing of its R1 reasoning model might also hurt adoption of closed-source models.

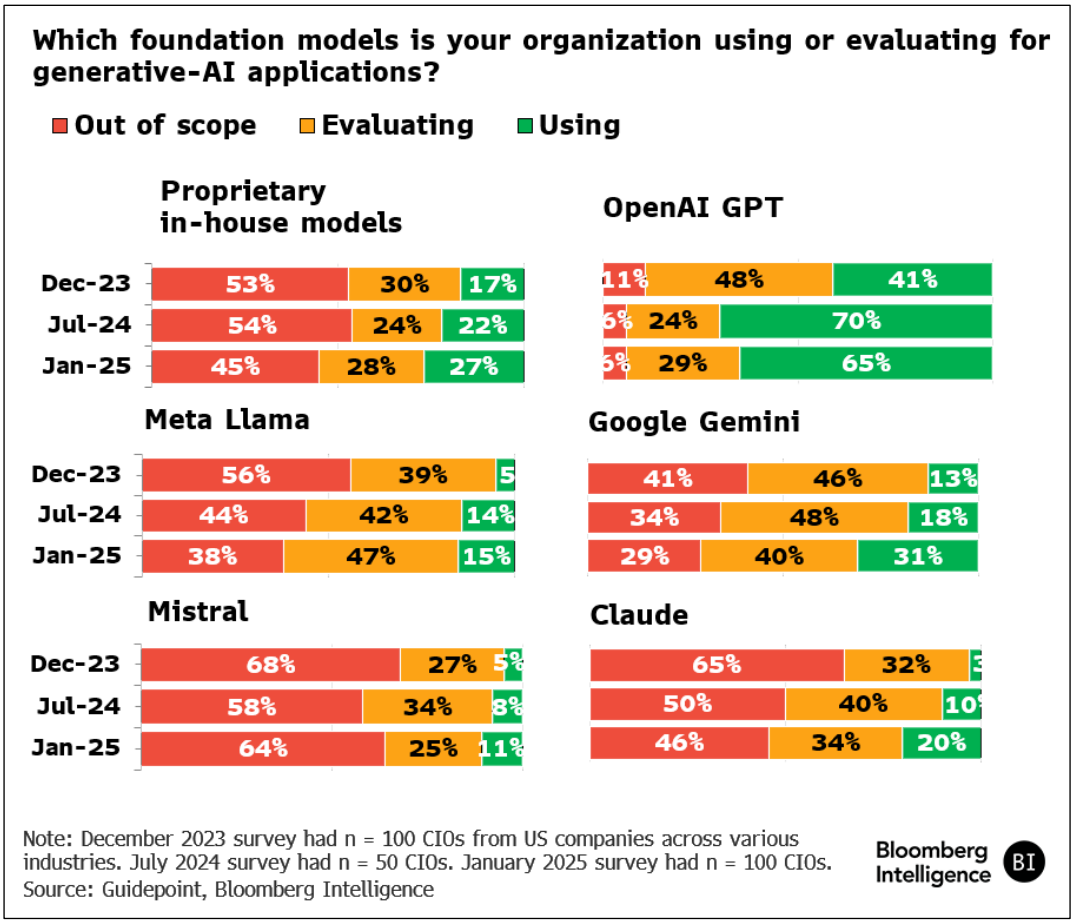
BI **DeepSeek disruption can temper use of closed-source models**

Participants in generative AI signaled a preference for Microsoft Azure cloud and OpenAI's GPT model over other foundational LLMs for inferencing workloads. The integration of Microsoft's Azure platform with OpenAI models continues to be an advantage over public cloud rivals for hosting inferencing workloads. Others, such as Google Cloud with its Vertex AI, offer application programming interface integration for all LLMs, including Anthropic's Claude, Mistral AI and Meta's Llama. OpenAI GPT isn't offered on Google Cloud or Amazon Web Services.

Microsoft's platform was cited by 28% of respondents for ease of deployment, while 17% named Google and 22% picked Amazon Web Services. Microsoft's lead is probably due to its early tie-up with OpenAI, which can help further narrow the gap with AWS in the cloud infrastructure services market. At the end of 2023 Microsoft had 16% of the market, compared with AWS at 47%. In 2018 Microsoft had 12%, while AWS had 48%. We expect the gap to close further by the end of 2025 as Azure's sales growth may outpace AWS' by 1.5x-2.5x.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 21: Foundational Model Preference**




Oracle and Salesforce are well-placed to gain cloud sales as enterprises look beyond hyperscalers on AI. Our latest CIO survey indicates that Oracle could become the No. 4 cloud infrastructure provider, after AWS, Microsoft and Google, while Salesforce taps its strong CRM market share.

Oracle Cloud Infrastructure's demand among survey respondents increased about 4 percentage points for AI training and remained the same for inferencing workloads. In January, 10% of responses indicated plans to increase spending on OCI for AI inferencing, compared with 11% in July's survey. For training, the figure rose to 11% from 7%, possibly aided by access to Nvidia GPUs, Oracle's competitive pricing and its ability to significantly increase capital spending.

BI **Oracle on path to be No. 4 cloud infrastructure provider**

CIO spending plans for Salesforce regarding AI training, inferencing and copilots logged among the largest increases outside the Big Three hyperscale cloud providers, suggesting that the company's new Data Cloud and adjacent Einstein One products are gaining momentum. Data Cloud enables customers to combine and optimize data for training and inferencing various generative AI models on the Salesforce CRM platform.

The increases in our survey also could indicate customers' preference for having their data aligned between the platform and the AI workload application, likely resulting in a more secure and better end-user experience.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

Our survey reveals a sustained rise in AI training workloads, driven by growing availability of data. This trend is likely to boost demand for data-center graphics processing units (GPUs) and custom AI training accelerators over the next 6-12 months. More than eight in 10 respondents expected capital spending on AI servers to increase in 2025, compared to only 32% for traditional servers.

Compared to AI inference tasks, model-training will require high-performance chips deployed in data centers that have stronger computational power,precision and larger memory capacity to handle larger datasets' parallel calculations.

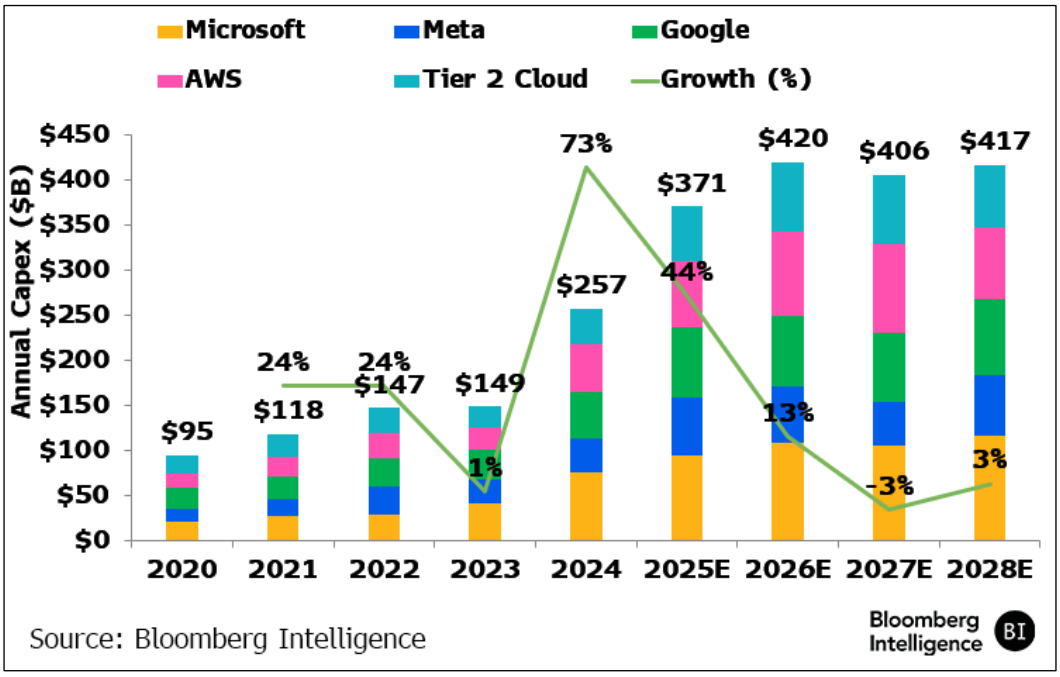
## 5.4 LLM Training

The continuous scaling of foundational LLMs will be a tailwind for both cloud hyperscalers, which have already seen a multi-billion sales contribution from AI workloads, and companies licensing their LLMs using APIs and broader enterprise use. We expect capital expenditures on LLM training to remain robust over the medium term given the high barriers to train a foundational model, which has already driven consolidation among LLM vendors.

Training of LLMs remains the largest portion of the Bloomberg Intelligence Gen AI forecast through 2032. With the five leading foundational model players – OpenAI, Google, Meta, Anthropic and Mistral – focused on scaling the next version of their LLM using the latest GPUs, we expect training costs to expand in proportion with the growth in LLM parameters. While small language model variants such as OpenAI GPT-4omini and Gemini and Llama lighter version have been released for on-device and edge deployment, the recent release of its chain-of-thought reasoning o1 model from OpenAI suggest the investments in scaling of frontier models is likely to continue in the near to medium-term. Training is expected to reach $580 billion, while inferencing could be around $735 billion market with the growth more back-end weighted for the latter, based on Bloomberg Intelligence market sizing model.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 22: Hyperscaler Capex**




BI

**Survey suggests sustained rise in AI-training workloads**

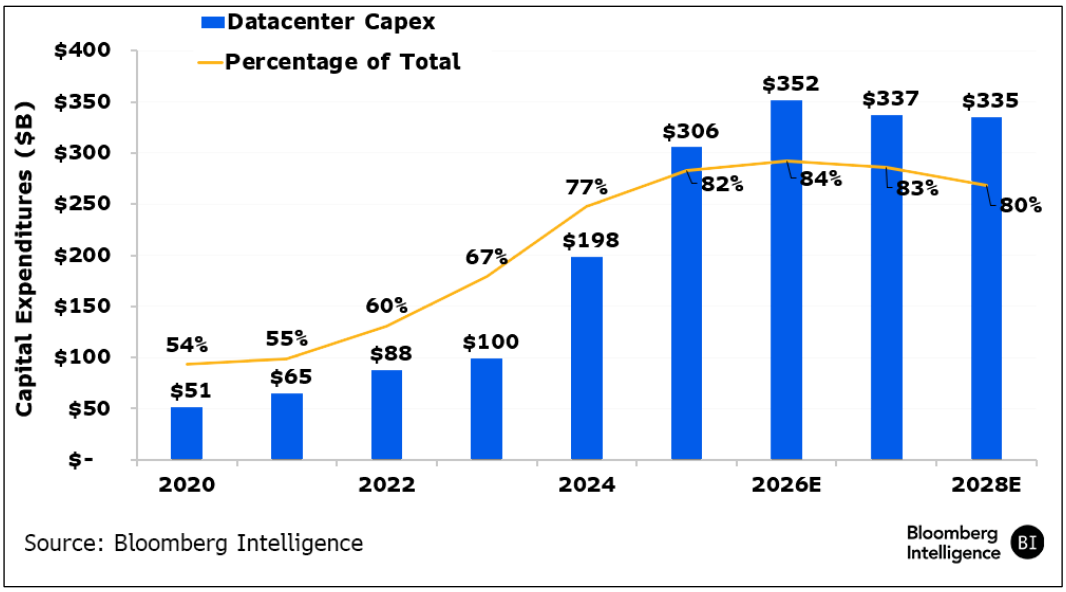
AI model-training costs should be similar for the leading LLM players including OpenAI, Anthropic, Meta, Google and Mistral, which all have a multimodal focus. Some of the image and video-specific LLM companies may have lower training costs, though we believe monetizing via APIs will involve leveraging the distribution of hyperscale cloud platforms.

Capital expenditures on data centers (see Fig. 23) are likely to stay robust for hyperscale cloud vendors including Microsoft, Amazon, Google and Oracle, which have already seen a multi-billion revenue contribution to their cloud infrastructure business from AI workloads. Microsoft's Azure AI business is already at a run-rate of $13 billion, about a high-teens contribution to its Azure segment sales of around $80 billion. Similarly, Amazon and Google have called out multi-billion sales contribution from AI infrastructure. Though new capex for data-center expansion is likely behind the above-70% hyperscale capex growth for 2024, we believe growth may taper gradually than abruptly given the faster architecture releases for GPUs and strong demand for both AI training and inferencing.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 23: Inside Datacenter Spending % of Total Capex**




BI

**AI data-center spending growth gradually declines**

GPT-4 was trained using a cluster of 25,000-30,000 A100 GPU chips, costing about $300 million for one training run, yet the newer models will likely require a bigger-sized cluster proportional to the growth in model parameters. Xai is building a 200,000 H100 cluster, while Nvidia CEO Jensen Huang has painted a vision for connecting one million GPUs in an AI factory. OpenAI expects to spend more than $200 billion through 2030, based on reporting from The Information, the majority of which will be spent on training its LLM.

SLMs require significantly fewer resources to train and fine-tune than LLMs, which could lessen capacity constraints seen by Microsoft and many hyperscale-cloud providers over the past year. In an example offered by Sebastien Bubeck, former vice president of gen-AI research at Microsoft, Phi-1 needed only 8 GPUs over the course of a week to train vs. an LLM that could require thousands. IBM reiterated this sentiment, stating that SLMs may be roughly 10% of the cost compared with LLMs. SLMs employ syntheticallycreated data in their training sets that facilitate learning without the need for massive amounts of input data.

Lower costs are also likely to trickle down to the end customer. In the case of Khan Academy, Microsoft's latest Phi-3-medium was able to produce similar results to ChatGPT-4 at almost no cost vs. the $30/month price of the latter.

The lightweight nature of SLMs also unlocks a wide range of industry-related use cases that require a foundational model to run on-device or in highly regulated sectors where data can't be sent to cloud servers. Additionally, given the lower requirements for training data, an SLM could be tuned to a far wider array of cases. Among industry verticals, we expect the financial, health-care and public sectors to be the most likely big adopters of SLMs given their unique regulatory constraints and sensitive data. Mobile applications or gaming companies could also benefit from SLMs as they prioritize the lower latency of an on-device foundational model.

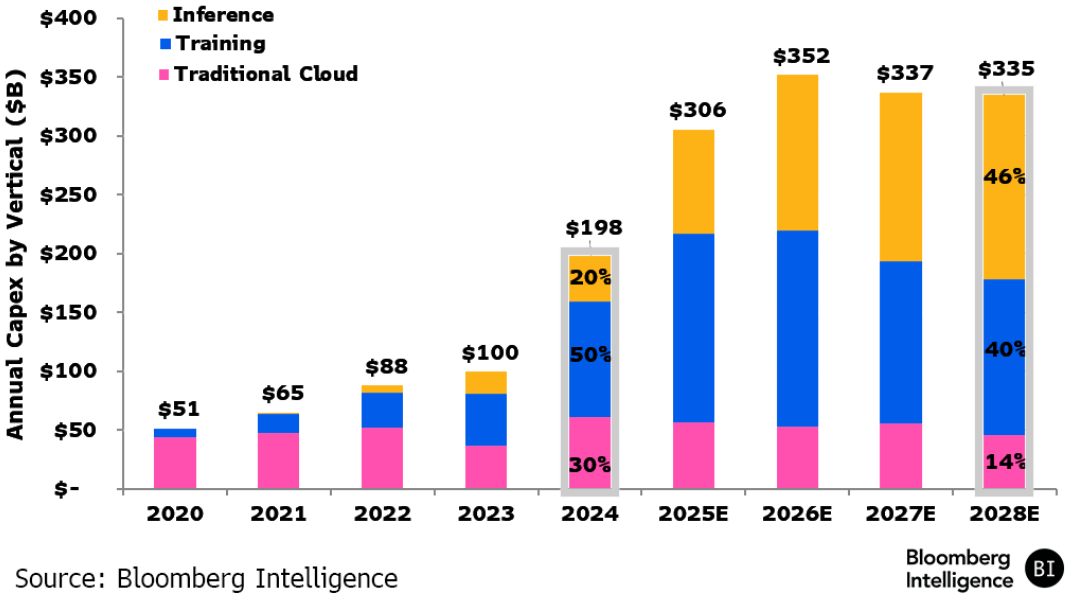


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

Meta is using its growing GPU compute clusters to power the training and deployment of its Llama LLM across its family of apps. Hyperscalers including Microsoft, Amazon, Google and Oracle have seen their cloud segments' growth accelerate, driven by AI-infrastructure workloads.

BI **AI spending can reach85% of all data-center capex in four years**

**Figure 24: Training vs Inferencing AI Spending**




Source: Bloomberg Intelligence

Bloomberg Intelligence BI

Recent talent acquisitions by hyperscalers including Google-Character AI, Microsoft-Inflection and Amazon-Adept AI suggest continued consolidation among the foundation model companies. Given the continuous scaling of LLMs and elevated capex requirements for data centers needed to deploy generative AI, we expect very high barriers to entry for companies looking to train their own foundational models. With the supply constraints around the latest GPUs, which have at least 2-3x performance and computation advantage over the prior generation, we believe hyperscalers are unlikely to pull back on data-center spending through 2026, barring a plateauing of LLM parameters or model performance.

Apple maintains the lowest capex intensity among big tech peers, partly due to the fact that it isn't a hyperscale cloud vendor that can leverage its distribution to support multiple LLMs. Capex intensity for Google and Meta could remain around 15-20% for the medium term, as these companies seek to scale the next version of their AI models. Though capex growth is likely to slow gradually over the next five years, a boost in cloud revenue could offset any headwinds to free cash flow. Pure-play gen AI LLM vendors like OpenAI, Anthropic and Mistral are likely to rely on licensing sales while partnering with hyperscale cloud providers to offset pressure from the elevated costs of AI training.

## 5.5 Inference Efficiency

Inference efficiency will likely become a big focus given the higher compute costs for OpenAI's o1 model, which the company is seeking to mitigate with an o1-mini version. Though the o1 model may offer an improvement in reasoning capabilities, both Meta's Llama and Mistral's latest large



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

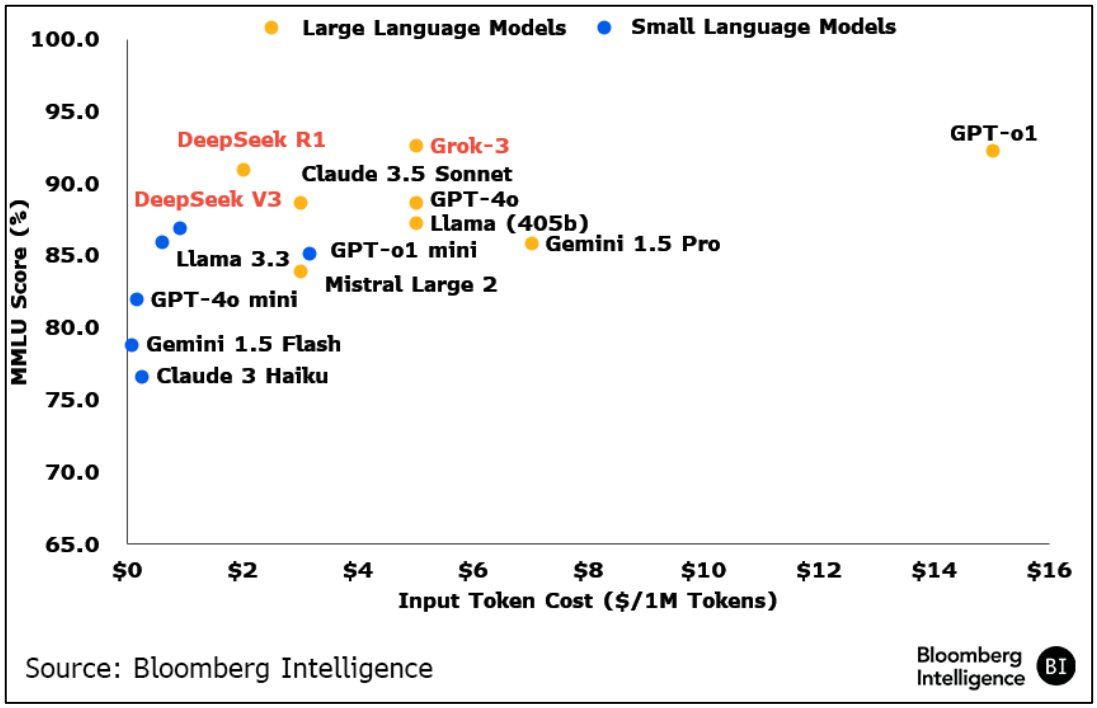
language model have incorporated multimodal reasoning and offer cost advantages for inferencing vs. OpenAI's models. Given the difference in pricing for different versions of the frontier LLMs, we believe companies are likely to mix LLMs from different providers depending on the type of query and to implement checks and guardrails by cross referencing each other LLMs.

**OpenAI's o1 may provide boost to Llama, Mistral**

OpenAI's release of its o1 model may further aid adoption of open-source models like Meta Llama and Mistral as enterprises increasingly focus on privacy, cost and fine-tuning LLMs with their own data. Given the likely convergence in model functionality between OpenAI GPT, Google Gemini, Meta Llama, Anthropic Claude and Mistral, we believe the lowest-cost model provider will be based on open source. OpenAI's o1 models performed better on certain benchmarks such as reasoning and math, suggesting different models may be employed depending on the query type. Meta also disclosed its Llama model family has reached 350 million model weight downloads on Hugging Face, with 20 million in just the past month. Access to Llama models via API at cloud partners grew 10x from January to July.

A pivot to smaller models with fewer parameters that can be used for specific tasks vs LLMs is likely to drive a secular shift in applications to agent functionality. OpenAI released its "chain-of-reasoning" o1 model, along with an o1-mini version, to showcase improvements in reasoning capabilities. Given the difference in pricing across different versions of LLMs, we believe companies are likely to mix LLMs from different providers depending on the query. A small, low latency AI model (5-10B parameters) will be included in iOS18 as part of the Apple Intelligence framework, that will be able to understand user commands, the current screen, and take actions on apps. It can handle tasks like summarization, as well as powering the "AI agent" features of Siri including user commands that require utilizing multiple apps.

**Figure 25: LLM Costs vs Performance**






No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

The mid-tier Sonnet model of Anthropic's Claude (an AI assistant) showed a notable performance improvement while lowering input and output token costs to below that of GPT-4o, suggesting competition is likely to increase around reducing AI-inferencing costs. This could give an advantage to hyperscalers such as Google, with its Vertex AI -- where it offers both its Gemini model and other foundational models through the Vertex AI platform. Amazon Bedrock also offers Anthropic's Claude models for inferencing workloads on its platform.

As LLMs grow exponentially larger, the number of floating point operations (FLOPs) generally scale with the parameter counts demanding significantly greater computational resources. To offset rising LLM costs, foundational model companies will seek to shrink the size of trained models for lowering inferencing costs for broader use cases across enterprise and consumer applications.

BI

**Greater need to reduce cost of AI inferencing**

Most foundational model providers, including OpenAI GPT, Anthropic Claude, Google Gemini, Meta Llama and Mistral, have released smaller versions using quantization and distillation for edge use cases to run LLMs natively on PCs and smartphone devices. We believe shrinking the parameter size of models will become more important amid continuous scaling of datasets and tokens used in the transformer architecture that underpins most foundational LLMs.

## 5.6 Open Source / Commoditization of LLMs

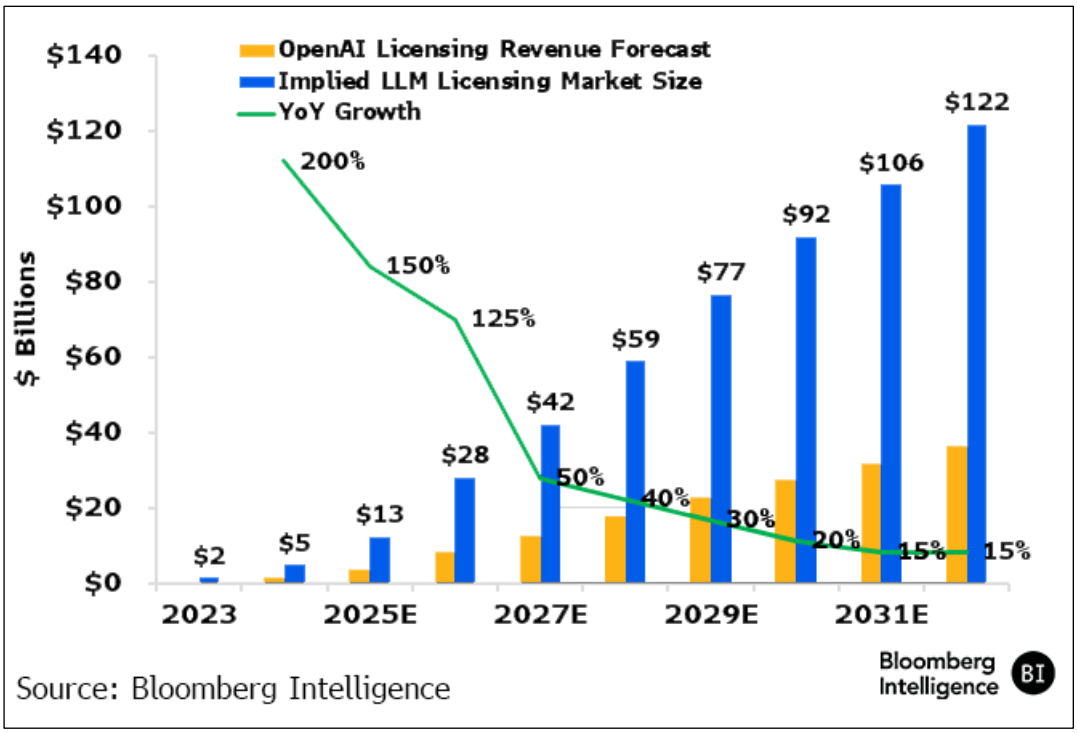
With most model providers having access to a similar corpus of web data, it will be hard for them to maintain differentiation based on dataset and compute. Model parameters are likely getting better with every new version, and inference costs could come down over time. OpenAI's o1-mini and o1 models are aimed at addressing different types of use cases involving chain-of-thought reasoning, though higher inference costs can be a deterrent.

Given the likely convergence in model functionality between OpenAI GPT, Google Gemini, Meta Llama, Anthropic Claude and Mistral, we believe the lowest-cost model provider will be based on open source. OpenAI's o1 models performed better on certain benchmarks such as reasoning and math, suggesting different models may be employed depending on the query type. Both Meta Llama and Mistral's latest open-source Pixtral 12B parameter model have closed the gap in multimodal functionality vs. closed-sourced LLMs like OpenAI's.

OpenAI recently raised $6.6 billion, implying a $157 billion valuation, putting a focus on the company’s monetization as pure-play generative AI player and licensing sales for its foundational LLM. Though training costs are relatively in the same range for most foundational AI model companies, greater monetization of OpenAI's technology and subscription products suggests room for Meta and Google to bolster their licensing API revenue, given their distribution advantage as they narrow the functionality gap vs. OpenAI.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 26: OpenAI Revenue Expectations**




BI **LLM licensing growth tapers, reaches $122 billion in 2032**

OpenAI's API revenue of $1 billion (30% of total sales) has been driven by its leading large language models and close integration with Microsoft Azure AI. OpenAI generates API revenue from both enterprise and individual application programming interface calls, as well as consumption through Azure AI. This is in contrast to Meta, which has open sourced its Llama foundational model and requires a license fee only for scale deployments. A narrowing of the functionality gap between the latest version of LLMs creates an opportunity for Meta and Google to similarly monetize their models via API and licensing revenue.

Hyperscaler distribution has so far aided Anthropic's API sales, at about a $500 million run rate, supported by the company's alliances with Amazon and Google. In contrast, OpenAI's licensing revenue from Azure may be about $200-$300 million, a mid-single-digit percentage of Microsoft's Azure AI workload sales. A lack of multi-cloud options beyond Microsoft Azure could limit OpenAI's indirect API sales, whereas other foundational models might see faster growth. Both Meta and Google have invested aggressively in building their own generative-AI infrastructures for model training and deployment. We expect API usage and other consumption-based contracts will be a big driver of LLM licensing sales in the near-to-medium term.

## 5.7 Gen-AI Coding Copilot Adoption Likely Above Other Uses

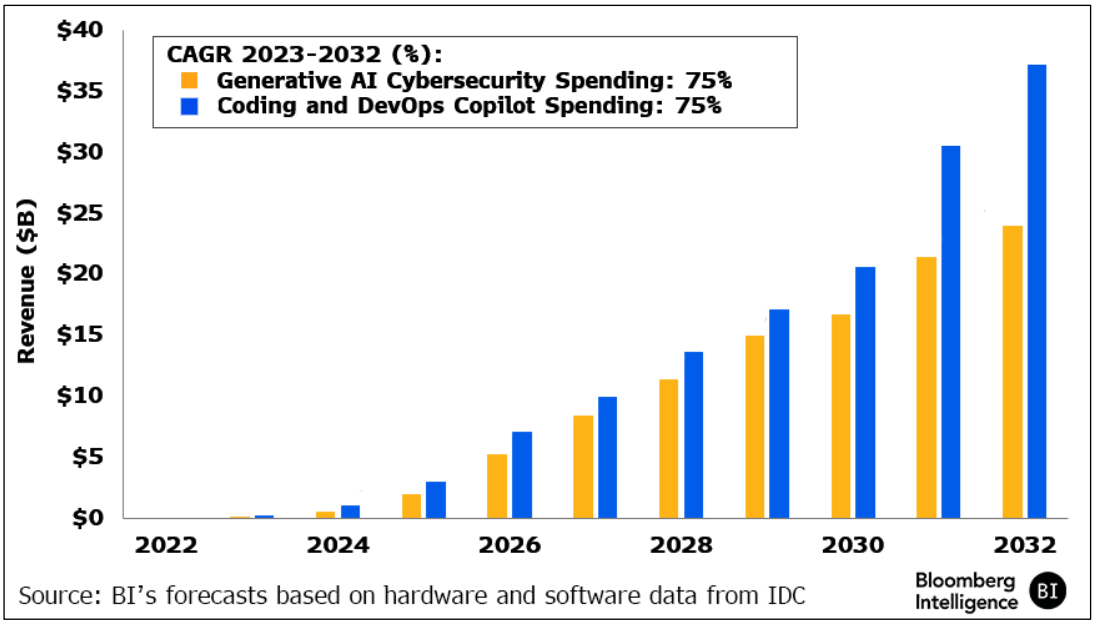
Coding and DevOps copilots may be one of the largest areas of generative-AI customer spending, and we anticipate adoption to be higher in these products than other gen-AI tools. These possess the ability to significantly shorten the time to a return on investment for customers, as they help increase a software engineer's productivity and ease pressures companies face from the global developer shortage. Tools such as Microsoft's GitHub Copilot and Amazon's Q



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

Developer have claimed over 40% greater coding efficiency. We expect that, even if half of that metric were realized by companies, it would be enough to drive outsized adoption.

**Figure 27: Generative AI Cybersecurity, DevOps**




BI **Coding, DevOps shorten the time for a customer's ROI**

GitHub is the clear leader at well over 100 million developers, and is is uniquely positioned as the biggest code repository for developers irrespective of the cloud provider, with annual recurring revenue potentially reaching $1.5-$2 billion by the end of 2025. Amazon's Q Developer and Google's Gemini Code Assist are also likely to expand within their developer bases, while growing integrations with development environments like Visual Studio could serve to broaden their compatibility and reach.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 28: Copilot Offerings**

| | GitHub Copilot Business | Amazon Q Developer |
|---|---|---|
| **Price** | **$19/Month/User** | **$19/Month/User** |
| **Productivity Boost** | **55% Faster Coding** | **Up to 40% Increased Productivity** |
| **Features** | • Code Completions<br>• IDE/Mobile Chat<br>• CLI Assistance<br>• Security Vulnerability Filter<br>• Code Referencing<br>• Public Code Filter<br>• IP Indemnity<br>• Enterprise-grade Security | • Code Completions<br>• Support for IDEs, Command Line<br>• Optimized for AWS Services<br>• AI Powered Code Remediation<br>• Reference Tracker<br>• Bias Avoidance |

Source: Bloomberg Intelligence

Bloomberg Intelligence BI

Greater adoption of these coding copilots -- which we expect to be more than 75% for GitHub and Q Developer over the long run -- could also accelerate the modernization of on-premise applications to the cloud, thereby also benefiting hyperscale cloud providers. An internal Amazon study noted that Q Developer upgraded over 1,000 Java applications in two days, a process that could take a team of developers roughly a year to complete, according to AWS Vice President Deepak Singh, who was featured on BI's Tech Disruptors podcast.

BI

**'Lift and shift' can speed sales in bid to modernize**

We anticipate that more than half of all cloud applications were brought over from on-premise without optimization, a process known as "lift and shift," which may also accelerate developer copilot sales as teams look to modernize these apps.

## 5.8 Data Comes at a Premium

The size and complexity of large language models makes the training process extraordinarily data intensive. Though OpenAI's ChatGPT reached a partnership with Microsoft, it still could be disadvantaged compared with internet enterprises in the volume of training data available.

ChatGPT's initial application was mainly focused on changing the nature of search, which has been dominated by Alphabet's Google. ChatGPT's primary use was to analyze, generate and edit text based on user input. Yet within just a few months, OpenAI realized how powerful generative AI can be, and the platform's scope quickly expanded beyond traditional search. The latest version of ChatGPT can handle data including images, audio and video. Such inputs require vastly more computing resources than text-based LLMs.

The size and complexity of LLMs based on transformers architecture are likely to grow as a result of multimodal input, which can help hyperscale companies including Microsoft-OpenAI, Meta,



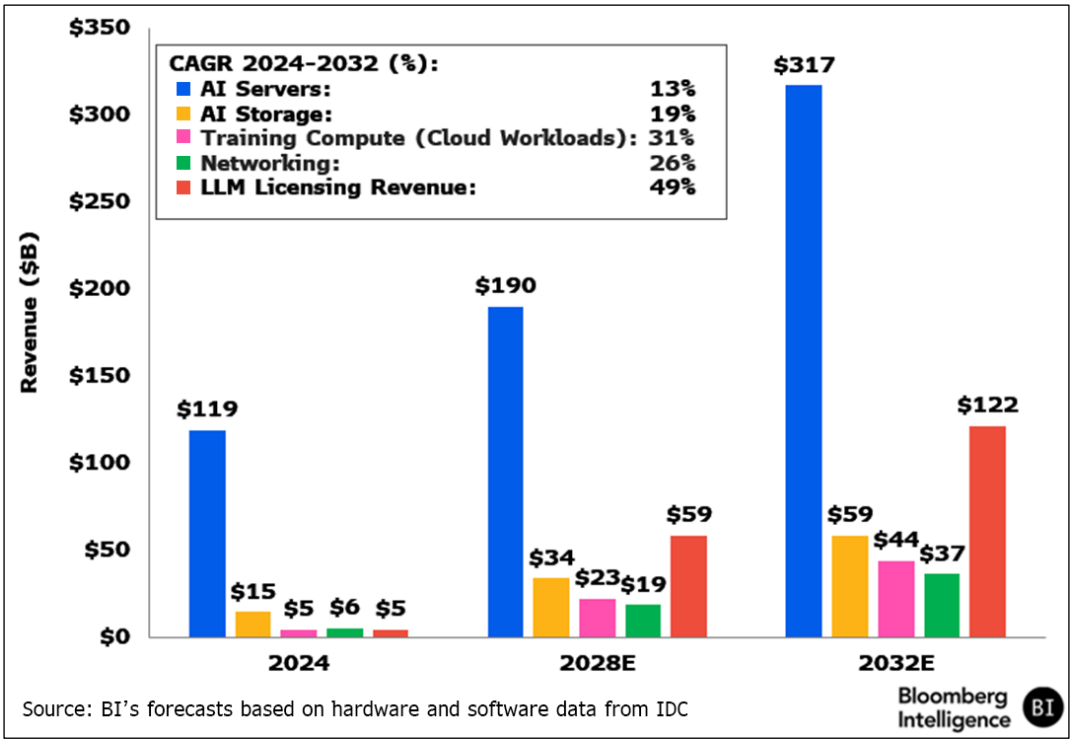
No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

Google and Amazon to maintain their lead over other foundational LLMs from peers. Given training of LLMs is a recurring process, we expect foundational LLM companies will consolidate to the ones that have a hyperscale infrastructure and large amounts of first-party data to improve the accuracy of the models.

Retrieval-augmented generation (RAG) techniques will likely become key for developing enterprise chatbots. This approach features RAG model that is trained on proprietary enterprise data, adding an element of customization. When a user inputs a query, the RAG model can identify relevant information from the enterprise database and append it to the user query, which then gets passed on to the LLM. The additional context provided by the RAG model helps reduce irrelevant answers and hallucinations from the LLM. Given the recent scrutiny around Alphabet Gemini's image generator, we anticipate hyperscalers will have a heightened focus on reducing historical inaccuracies and questionable responses. As a result, these companies could be first in line to develop a RAG counterpart to their existing models.

BI **Training compute market grows at nearly 31% CAGR through 2032**

**Figure 29: Training Forecasts, 2023-32**




## 5.9 AI Fuels a Fifth of Global Server Revenue

Robust gains in the number of ChatGPT active users indicate that generative AI could be among the most important catalysts to growth for the server supply chain in coming years, driving over 20% of global server revenue by 2024 from 15% in 2021, by our calculations.

After its November launch, OpenAI's ChatGPT amassed a base of 1 million users in a week and exceeded 100 million in just two months. OpenAI introduced a subscription service at $20 a



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

month and offered businesses paid access to ChatGPT to expand commercial applications. Companies including Snap, Shopify and Instacart already have integrated ChatGPT into their products.

The server supply chain's original design manufacturers could reap the most demand, since cloud service providers are integrally involved in AI development. AI servers could also drive robust sales for other suppliers with design expertise.

Fine-tuning of large language models, application modernization and data-related services are a few ways that IT services companies can leverage increased generative-AI spending. We anticipate consulting-heavy firms like Accenture, IBM and Capgemini to see more of the outlay than offshore providers like Infosys and Wipro.

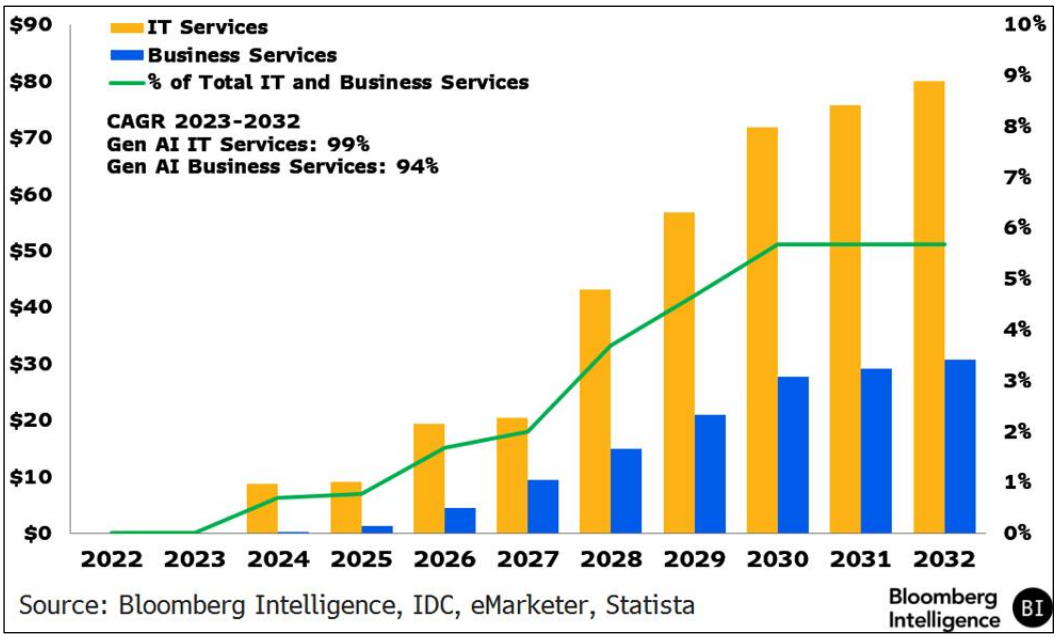
The services industry will likely play a major part in fine-tuning of large language models (LLMs), as end markets look to maximize the benefits provided by copilots. We see this work as akin to customization of packaged software, which has been a key backbone of the industry for the past 20-30 years. Data-related services is another area that can play a big role, which includes data aggregation, cleanup, creating common data lakes and using that final information to create custom copilots for clients.

Application modernization could also see increased use of gen-AI technology. The IT services industry will likely lag behind AI infrastructure and software, like the trends observed in past emerging technologies. Our analysis indicates meaningful growth starting in 2026 onward.

The total market for IT and business services is valued at $1.2 trillion and could reach $1.9 trillion by 2032, assuming 5% yearly growth. This implies generative AI could eventually represent 5.7% of the market by then, which we think is a conservative estimate.

BI **Gen AI can add $80 billion in incremental IT services sales by 2032**

**Figure 30: Generative AI for IT & Business Services ($ Billion)**




The shift to cloud and embracing more digital workflows is another area which will likely see



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

increased adoption because of gen AI, especially as majority of the IT spending still resides on-premise, where data is siloed across multiple systems. The recent success of Salesforce's Data Cloud product highlights the need for cleaner data to train LLMs. Cloud is likely to be another essential precursor to AI workloads, which could result in an accelerating shift from on-premise architectures. We see the services industry playing a big role in this effort.

Accenture, IBM and Capgemini are well-positioned to capture more generative-AI projects than the India-based IT services companies due to their larger consulting units. All three have invested in emerging technologies, with Accenture committing $3 billion to expand its data and AI practice, Capgemini making a similar €2 billion investment and IBM enhancing its software portfolio, which could drive greater work for its consulting unit. These investments position the companies well for when corporate IT spending improves, especially compared with peers with less robust gen-AI offerings.

Services spending could increase about $113 billion over the decade, doubling annually from 2023. Consulting, data-related services, custom application development and creation of new chatbots would drive the additional expenditures. The current total market for IT and business services combined is roughly $1.2 trillion, which may reach $2.1 trillion in 10 years, assuming 6% annual growth.

Among IT services peers, Accenture stands out given its propensity to invest ahead of the curve in emerging technologies. In 2023, the company announced a three-year $3 billion investment in generative AI that would include training 80,000 workers on AI as well as offering greater solutions and models. We calculate generative AI will contribute roughly $1.5-$2 billion to Accenture's revenue by 2025, after booking $450 million alone in 1Q. Outside of Accenture, practices such as those from IBM, Infosys and Tata Consultancy Services will also likely benefit.

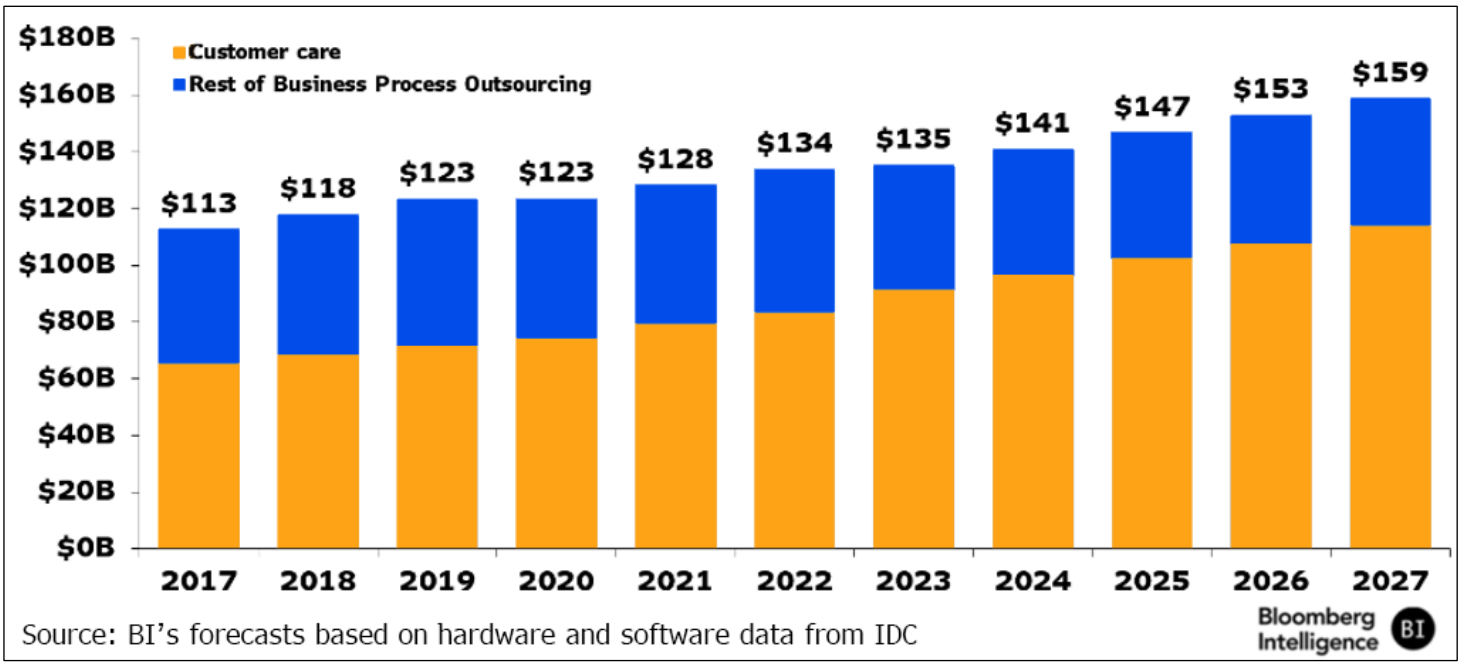
BI **Accenture on a path toward $2 billion in AI revenue**

Business-process outsourcing services (see Fig. 31) may be more heavily disrupted than IT services, with jobs in customer service and back-office areas displaced by AI assistants. That could lead to near-term pricing pressures, particularly in customer care, which sits at the bottom of the BPO value chain yet is still the largest and fastest growing of its segments. Customer care is forecast to expand 6.5% annually through 2027, compared with 3.5% for all other BPO services, according to IDC.

Generative AI will be more of a revenue tailwind for BPO companies that have minimal exposure to customer-care services, like Genpact and EXL Service. Companies that specialize in greater value-added services might leverage efficiency gains from AI to expand their total addressable markets, particularly in areas like data analytics.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 31: Business Process Outsourcing Forecasts**






No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Section 6. Expanding Uses

## Enterprise Gen-AI Use Cases Rising, Adding Productivity

Generative AI has the potential to boost productivity across multiple areas, including software development, customer service, operations and security. As language models evolve, costs are likely to drop and enterprise use cases expand. Though much of the quoted AI success is from companies developing these tools, more customer testimonials might drive greater adoption.

**Figure 32: Generative AI Use Case Tracker**

| End Market Impacted | Quote | Company | Area of Company |
|---|---|---|---|
| Software Development | With Q's code transformation capabilities, Amazon has migrated over 30,000 Java JDK applications in a few months, saving the company $260 million and 4,500 developer years compared to what it would have otherwise cost. | Amazon | Back End Operations |
| Software Development | In our engineering organization, our developers now save more than 20,000 hours of coding each month through the use of our AI tools. These innovations are helping us drive continued efficiencies across the business and accelerate our product road maps. | Salesforce | Back End Operations |
| Software Development | We've done an awful lot to digitize many parts of our business and we're now applying Gen AI to it. The places that we're seeing tremendous -- tremendous results on are coding. We need 70% less coders from third parties to code as the AI handles most of the coding, the human only needs to look at the final 30% to validate it. That's a big savings for the company moving forward. | BP | Back End Operations |
| Security | implementing and managing a scalable platform powered by Gen AI, enabling the agency to act on evolving cyber threats up to 60% quicker than with traditional technologies, including detection, response, and containment models. | Accenture | Back End Operations |
| Security | Initial modelling shows AI enhancements boost fraud detection rates on average by 20% and as high as 300% in some instances. | Mastercard | Back End Operations |
| Sales | AI-powered personalization can increase revenue per passenger by 10 to 15% | Microsoft | Front End Operations |
| Sales | Typically our sales cycle was 90 days and now with these experiences that HubSpot AI has helped us to build that, has gone down to 30 days to 45 days. We had people spending a minute or less on our experiences and our pages, now they are spending 3 minutes to 4 minutes. | HubSpot | Front End Operations |
| Sales | 81% of sales teams are either experimenting with or have fully implemented AI. The results speak for themselves: 83% of sales teams with AI saw revenue growth this year vs. 66% without AI. Sales Teams Using AI 1.3x More Likely to See Revenue Increase | Salesforce | Front End Operations |
| Operations | Nurses are spending, on average, 20 minutes an hour on admin tasks that we now with generative AI can help reduce to five minutes. | Koninklijke Philips | Back End Operations |
| Operations | While on the Lexis Nexus, what they're able to do on They built an AI platform, not just for conversational search, but also for summarization and intelligent legal drafting, which enables lawyers to be more efficient and be productive. They report their legal professionals are saving up to 11 hours per week using their AI platforms | LexisNexis | Back End Operations |
| Operations | AI has helped us in our fight against financial crime by reducing the processing time required to analyze billions of transactions across millions of accounts from several weeks to a few days. | HSBC | Back End Operations |
| Marketing and Advertising | Ad campaigns using Meta's generative AI ad features resulted in 7.6% higher conversion rate...More than 1 million advertisers and 15 million ads created with our generative AI ad tools, in the last month | Meta | Front End Operations |
| Marketing and Advertising | We use AI and tech to boost our agility on advertising and promotional expenses. You're familiar with this BetIQ program that we have started to help us allocate with the best ROI all our A&P which gives a good improvement of 10% to 15% to the productivity of our A&P and we're just beginning to roll it out. It should reach around 60% of our A&P at the end of this year. | L'Oreal | Back End Operations |
| Customer Service | They recently rolled out service now assist, generative AI product for their customer service piece. And they are now solving customer queries 55% faster because they're using now assist to help with things like case summaries, complex case notes. | BT Group | Front End Operations |
| Customer Service | Phone Pe leveraged Freshworks customer service platform to automate common service inquiries. Resolving 80% of inquiries without human intervention | Phone Pe | Front End Operations |
| Customer Service | Ada helped BlueJeans serve 72% of its Chat Users with a Chatbot vs. Live Agent. 75% of support contacts engage with chatbot. 83% reduction in live chats from trial end users | BlueJeans | Front End Operations |

Source: Bloomberg Intelligence

Bloomberg Intelligence BI



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

## 6.1 Productivity, Costs, Efficiency All Boosted

AI offers significant productivity benefits in software development, and coding is one area where we see generative-AI technology producing benefits for users (Fig. 33 shows some uses). Amazon Web Services reduced Java application time to a few hours from days with its AI-coding tool Amazon Q Developer. Similarly, Palo Alto reported a 30-40% productivity boost for developers using Copilot tools vs. those without. Companies like BP have seen substantial cost savings, needing 70% fewer third-party coders due to AI handling much of the work.

Customer service is another area well-suited for generative AI, with potential for greater adoption across call centers and support teams, as it improves efficiency and shortens response times. Ada helped BlueJeans automate 72% of its chat interactions using chatbots, resulting in an 83% decrease in live chats during a trial period, and BT Group saw a 55% faster resolution of customer queries after implementing ServiceNow's Now Assist, an AI tool that streamlines tasks like case summaries and handling complex notes. Discover also found that AI lets human agents retrieve information 70% faster than traditional search tools, significantly boosting agent productivity and customer satisfaction.

Generative AI is also having a significant effect on operations and back-office functions, driving productivity and reducing time spent on routine tasks. Take the example of Royal Philips, which found that nurses who typically spend 20 minutes each hour on administrative duties, can cut that to just five minutes with AI. Germany's largest health insurer uses AI to process claims by reviewing 800 policy documents, reducing the time required to just three seconds from 23-30 minutes, according to Google.

BI

**Cost savings come in a range of ways**

Similarly, Mastercard reported that after only 28 days of using Microsoft Copilot, it had saved a total of 1,200 hours, or about four hours per employee, in three specific use cases.

AI is making significant strides in improving product capabilities and bolstering security measures across various sectors. HSBC has leveraged AI to drastically cut down the time required to analyze billions of transactions across millions of accounts, reducing the process from several weeks to just a few days, while Mastercard has reported that its early AI implementations have boosted fraud detection by an average of 20%, with some instances showing improvements of up to 300%.

These developments highlight how AI enhances the efficiency and effectiveness of financial services, letting companies proactively combat threats.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 33: Companies Using Generative AI**

| | |
|---|---|
| **Kroger** | Deliver the appropriate promotional offers and discounts to customers at the right time |
| **Target** | Power product detail pages and provide more friendly and relevant explanations |
| **T-Mobile** | Spends around $2.5 billion on advertising annually, using AI to place ads and optimize media spend |
| **eBay** | Enhance sellers' product images to make them more compelling to customers |
| **State Street Corp** | Digitize and automate approximately 85% of bank loans settlements |
| **Edison International** | Improve inspections, customer experience, and grid planning; Also using AI for research, workflow automations, and code development |
| **Hershey** | Optimize the talent profile for certain key jobs |
| **McKesson Corp** | Improve patient intake and workflow, system productivity, and several supply chain use cases, including disruption predictions, forecast accuracy algorithms and fraud detection |

Source: Bloomberg Intelligence

Bloomberg Intelligence BI

BI **Cross-sector applications for Gen AI**

## 6.2 Base Capabilities Cover Broad Range of Uses

Momentum with gen AI agents is set to accelerate amid successful proof of concepts and ideation across different types of consumer and business apps. Real-time language translation has helped boost engagement and satisfaction, and LLM-based recommendations will be connected to enterprise databases to automate certain business processes across websites, mobile apps and any other touch points. Newsweek is already seeing faster content discovery with its LLM-based search, with longer context windows than what was offered by traditional search.

We expect a broader change in user interface for all types of applications and software as free-form search is likely to play a bigger role replacing the traditional menu- and directory-based navigation. In addition, the use of AI for content ranking and recommendation is likely to drive more personalization across most platforms.

Developer coding copilots, text summarization and customer-service chatbots are among the most prominent use cases for generative AI (Fig. 34), and generating marketing content for websites and presentations with LLM-based prompts promises to be a big productivity driver. The biggest proofpoint is in the ad-targeting improvements that Meta has seen across its advertiser base, driven by its Advantage+ suite that leverages generative AI. Others have done the same, such as Alphabetith its Performance Max offering. Marketing copilots are already deployed at a number of enterprises for drafting documents or generating images based on prompts, which will likely help drive further automation for targeted promotions. Increased model efficiency and hardware improvements could enable more use cases on the consumer side, with



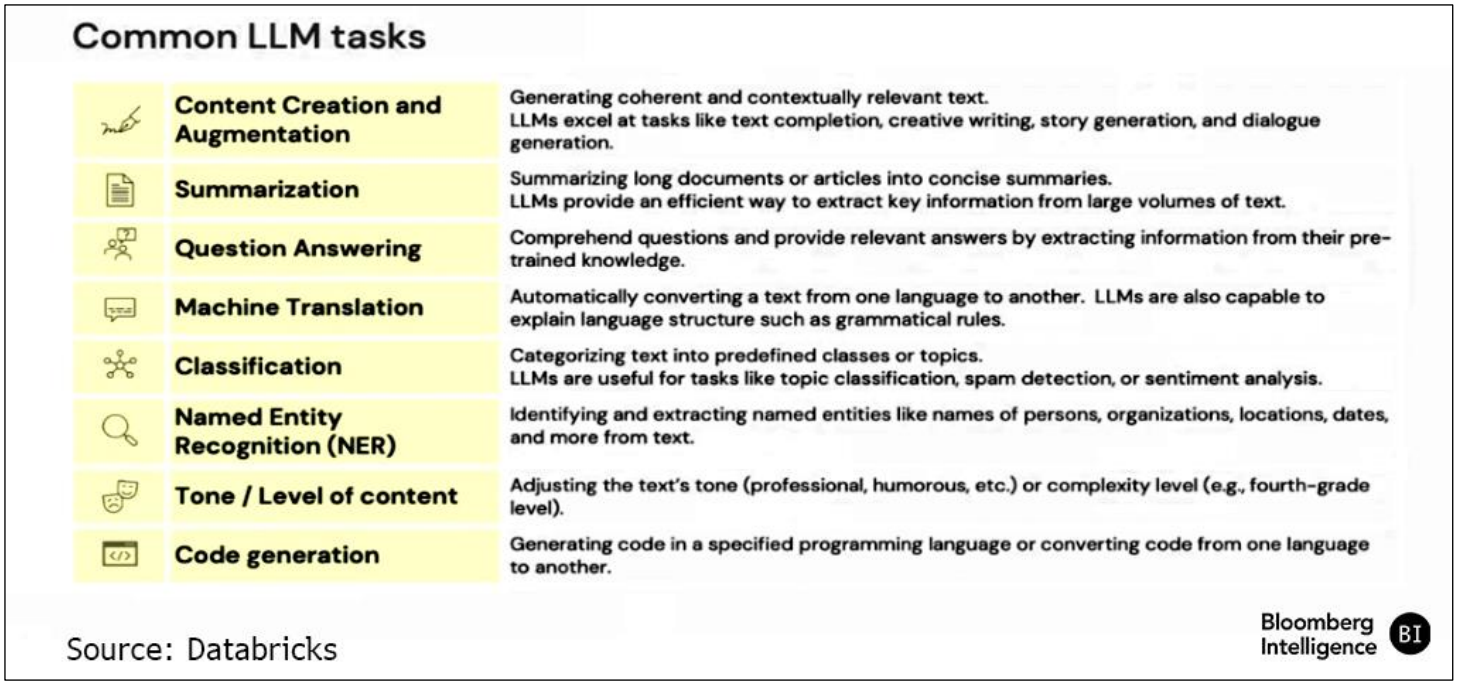
No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

apps likely to deploy user interface changes based on voice and gen-AI assistants to bolster engagement.

Amazon Web Services, Microsoft, Google and Oracle could win increased orders for cloud services, in addition to hybrid cloud providers such as IBM, for which the main source of traction would be in its Red Hat, security services and Watson-related products. Oracle can parlay its leading market share in database-management products, while Cisco, Databricks, Snowflake, VMware and ServiceNow also stand out. The leading cybersecurity players, including CrowdStrike and Microsoft, have launched security copilots.

Similarly, design and gaming companies including Adobe, Unity, Roblox and others are integrating AI into their software to fend off competition from startups that use LLMs.

**Figure 34: LLM Uses**




## 6.3 Apple, Amazon in the Mix as AI-Enabled Tools Expand

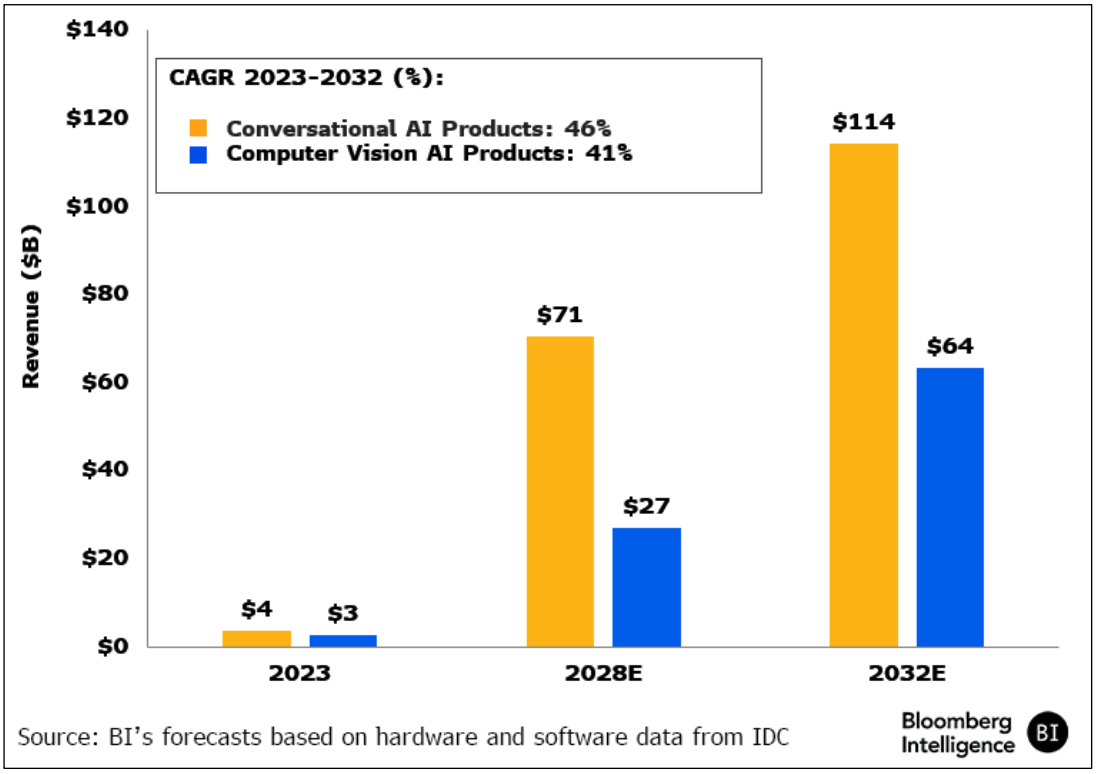
Voice assistants based on conversational AI and computer-vision products may emerge as new categories for inference, given the availability of large language models (LLMs) for domain-specific predictions. Apple, Samsung, Amazon and others may look to conversational AI given how well their existing product offerings mesh with the category. The launch of more compact versions of existing LLMs bodes well for consumer devices given the reduced computational intensity required to run AI workloads. Auto manufacturers like Tesla and GM could invest in computer vision research to drive the next generation of AI in vehicles. Advancements in generative AI and more accurate responses for recently trained LLMs have set the stage for such categories to accelerate the overall $1 trillion devices market, where smart speakers and wearables already are large categories. Mainstream adoption of generative AI could drive a faster


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

refresh cycle for personal computers and smartphones as new versions of these edge devices may be optimized to run generative AI apps natively, likely leveraging a smaller LLM given the processing, memory, and storage requirements. Adoption in Asia, albeit gradual, is being led by technology giants there. In one instance, Alibaba's open-sourced SeaLLM, despite having fewer parameters than OpenAI's ChatGPT-3.5, can process text in Southeast Asia faster and with more accuracy than the latter, thanks to a custom-built training dataset tailored for the region's diverse language profiles and cultural norms.

Retail is ahead of most sectors in AI investment, representing about 13% of all such spending in 2023, Statista data show, adding capabilities that can aid revenue and margin. We believe that generative-AI spending will increase as retail leans further into driving innovation and speed across search, marketing, product design, supply chains, customer service, seller tools and virtual-shopping assistants. Companies that use AI recorded a 700-bp increase in sales in 2022-23, doubling their profit pace, according to Statista, and its growing use is aimed at boosting conversion -- about 2.5-3% on average, according to Shopify -- and improving productivity and margins. About two-thirds of retailers aim to add Gen-AI features for their customers, according to eMarketer.

**Figure 35: Inference Forecasts 2022-32**




BI **On-device inferencing could generate almost $180 billion in revenue**

## 6.4 Getting Personal to Drive Sales

Continued integration of AI and notably generative AI can further strengthen a brand's appeal as it improves marketing messages to better meet and match customer expectations. Tools like augmented reality and virtual reality, tailored recommendations and personalization can all boost


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

conversion, while other companies are using generative AI to quickly formulate marketing campaigns.

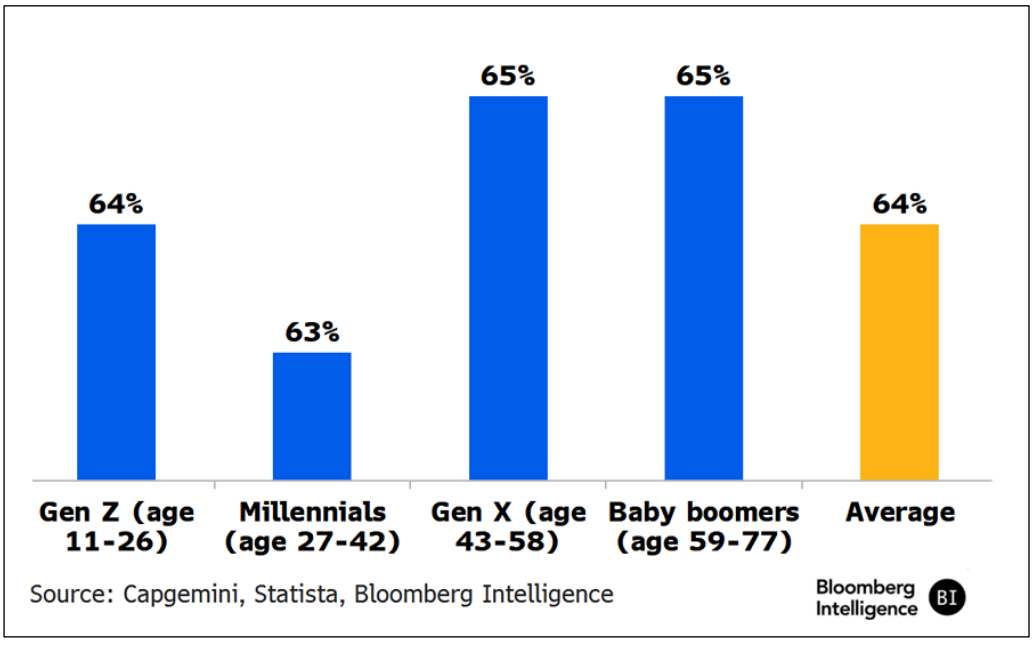
Offering AI-driven recommendations and similar-item features can boost conversion as it helps consumers find what they're seeking, and from Gen Z to Baby Boomers (see Fig. 36), consumers are comfortable with buying based on AI-based recommendations, with about two-thirds becoming open to shopping AI-generated suggestions, Statista data show.

Wayfair's Decorify and Muse leverage augmented and virtual reality to recommend items for a specific space and visualize items there, while Warby Parker uses AI to suggest frames based on a shopper's preferences and can show how them how they'll look on their face. Sephora and Ulta both use color match and virtual try-on technology to help shoppers find the right colors for their skin tones.

Amplifying the allure of product photographs with AI is becoming more common across retail, as it helps boost conversion and save costs. Several online marketplaces are using design features to enhance shopping experiences. Amazon, eBay, Wayfair and Google, among others, have layered in product-imagery tools to save time and money, and boost aesthetics. Revolve has used gen-AI to create its billboard campaign "Best Trip" as well as its capsule collection, which could pare marketing and design costs. Puma is also using AI to create personalized product images, such as customized backgrounds, to improve its marketing campaigns, while Under Armour used AI to create an inspirational "team talk" speech to encourage athletic activity and draw attention to the brand.

BI **Almost two-thirds of consumers say they're willing to use AI suggestions to buy products**

**Figure 36: % of Consumers Open to Buying AI-Suggested Product**




## 6.5 Conversational AI to Expand Voice-Assistant Category

Conversational AI products from hardware makers like Apple and Samsung will likely be tethered to PCs and smartphones, helping to boost upgrades of existing installed bases while driving


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

growth for services. Suppliers like Apple (HomePod), Google (Home) and Amazon (Echo) may enhance their device speakers with assistants, while carmakers including Tesla, BMW, Ford and Volkswagen could incorporate them to boost driver engagement. Conversational AI is much more popular with consumers than generative AI for copilots, according to a recent Bloomberg Intelligence survey, with over 40% of respondents citing frequent use of AI tools for conversational interfaces. We expect these products to grow at roughly a 54% compound annual rate through 2032, in line with the overall generative AI market. Most of the gains will probably come in the latter half of the period as the product category becomes more established.

Conversational generative-AI products could grow at a compound annual rate of 46% to $110 billion in revenue by 2032. Similar to the Google Assistant on Pixel 8, Apple previewed a revamped Siri powered by Apple Intelligence at its developers conference in June. It also partnered with OpenAI to leverage GPT for more advanced AI capabilities. Apple HomePod, Google Nest and Amazon Alexa could also see similar enhancements to their built-in conversational assistants.

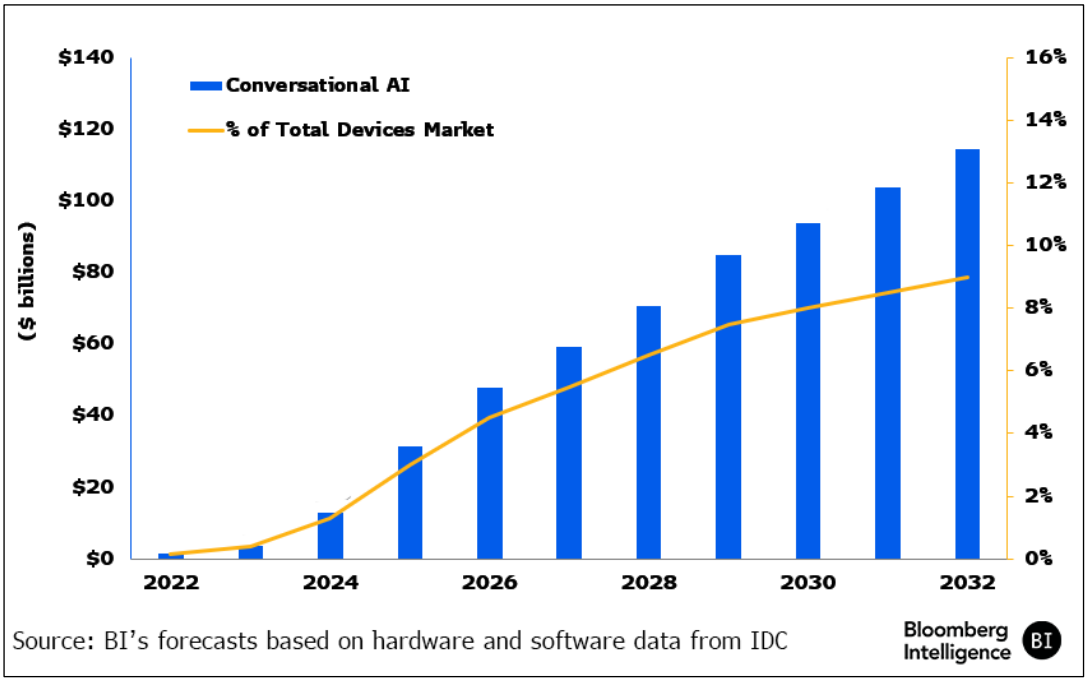
BI

**About four in 10 survey respondents use conversational AI frequently**

Conversational AI will likely affect many markets. Automakers might implement navigational assistants and companies could use it for HR training and health-care providers may design virtual assistants for patients.

Computer vision also could become a significant application of generative AI tools. Building LLMs will require large amounts of training data and then need generative AI for deployment in automobiles to run inference functions. We expect most incremental revenue from computer vision to come from hardware, with the category expanding to about $60 billion by 2032, thanks to its application in advanced driver-assistance systems. There may be an even larger impact on related service sales in the medium-to-long term.

**Figure 37: Conversational AI**






No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

## 6.6 Better Computer Vision May Boost Autonomous Driving

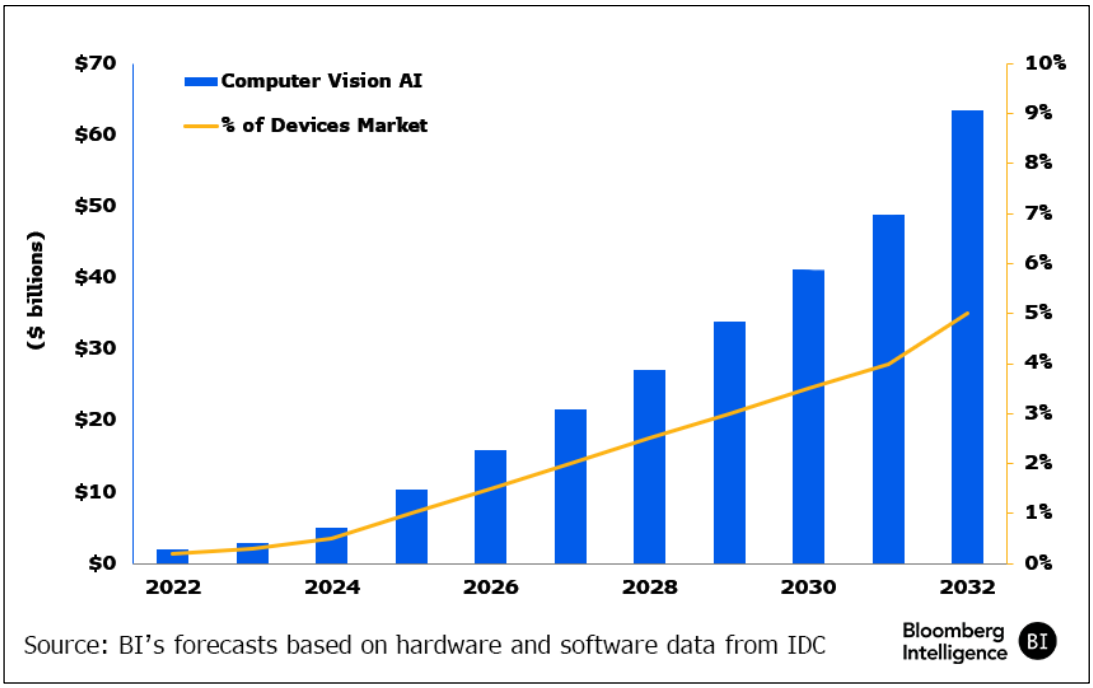
Computer vision is another major application of generative AI and could reach $58 billion in revenue by 2032. Autonomous driving may be a catalyst for computer-vision technology, particularly as Tesla aims to launch its robotaxi later this year. We expect multimodal AI's image- and video-processing capabilities to improve, which can be useful for applications like autonomous driving and medical scanning, given the need for accuracy.

Virtual reality (VR) might also benefit from improved computer vision, though VR-headset adoption has struggled amid discomfort with the form and pricing. Augmented reality could see better traction, aided by generative AI.

An AI training infrastructure will be essential to run these heavy workloads, sparking demand for high-capacity servers and storage. Most training-related workloads will be new since enterprises currently use general-purpose CPUs for analytics and transactions.

BI **Image, video processing aids industries from autos to medical devices**

**Figure 38: Computer Vision AI**




Generative AI can also improve the seller experience on marketplaces, with eBay, Amazon and Depop among those adding features that speed the listing process. AI tools are also making image listings cleaner, which can improve conversion, and AI authentication can also improve a platform's inventory.

BI **E-commerce is fertile ground for AI**

Generative AI can help sellers on e-commerce marketplaces, improving the quality of item descriptions and aiding inventory, as seen by. EBay's magical listing feature, including its bulk listing tool, uses generative AI to write item descriptions, titles and categories, and has a more than 90% acceptance rate. EBay also has AI-powered apparel authentication through its Certilogo acquisition, which can enhance its inventory assortment, and has integrated generative AI into its parts and accessories app to help shoppers more easily find auto parts.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

Etsy's Depop has a generative AI feature that auto-populates listings descriptions and attributes using a simple photo upload.

Amazon's third-party sellers can use AI tools to boost marketplace appeal and listings, bringing efficiencies that may draw even more sellers. The AI assistant Amelia answers sellers' questions, gives access to critical business metrics and cuts time to create and manage listings. Sellers can also use AI to quickly transfer listings to Amazon from their own websites. Enhancing the tool for bulk listings could widen product breadth.

Amazon's AI-generated A+ Content tool can help sellers create more engaging and informative listings, attracting shoppers and saving time and resources previously allocated to photo shoots, drafting and testing. Amazon says it can increase sales by up to 20%. The Video Generator tool can create ads for sellers by using an image to create a custom video.

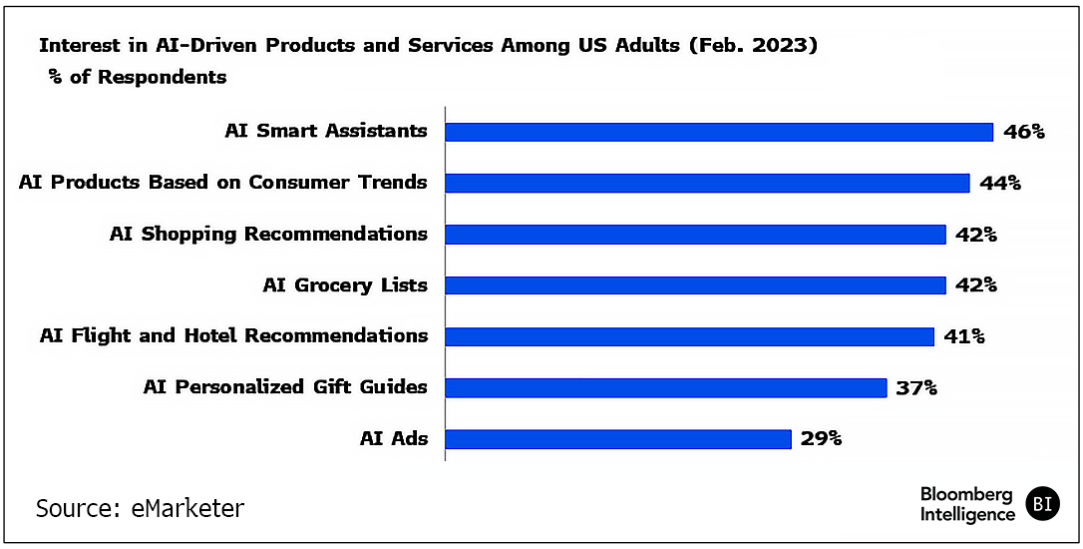
Generative AI is paving the way for chatbots to become personalized shopping assistants that take consumer requests and display appropriate brands and products. Snap, Meta, Pinterest and other companies that already have shopping on their platforms are investing in AI chatbots and could implement personalized shopping assistants to boost user adoption of social commerce, fueling monetization opportunities. Enhanced capabilities to handle longer and more flexible search prompts may also benefit retailers. Walmart demonstrated a search-enhancement feature at the Consumer Electronics Show – or CES - in January 2024, which accepts less-specific user input and fetches relevant products. For example, the prompt "help me plan a football watch party" resulted in products like chips, salsa and other game-day essentials being shown. We expect more retailers to add longer search-prompt capabilities to their apps or websites, which may significantly expand the number of search queries on all kinds of digital platforms, especially in e-commerce.

Multimodal search could enhance the user experience beyond text-based functions, which currently dominate the market. We believe the conversational nature of ChatGPT might reduce ad loads in the near term, as summarized responses reduce the need to click on links to find information.

**Figure 39: Service Interest**

BI

**Consumers like the idea of 'smart' assistants**






No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

As generative AI and machine-learning algorithms develop and adjust from user input, they cater to the person's tastes, interests and lifestyle, providing a more customized, unique experience and curating new content for social media and search platforms. That can expand availability and engagement, similar to how TikTok uses AI algorithms to recommend content to users.

LLMs may improve ad targeting for large companies that are rich with first-party data. Meta has already pivoted to AI-based recommendations with its Reels offering, helping to offset some of the headwinds from Apple's changes to its identifier for advertisers (IDFA) policy. Meta can continue to develop its Llama LLM and enhance the quality of its advertising campaigns.

Generative AI could also accelerate the shift to digital ads from linear TV, especially since offering personalized versions of advertising can increase efficiency and sales conversion. LLMs should also provide added benefit for existing large media companies as more premium content shifts to streaming from linear TV. Our analysis suggests that the generative AI market may add around $207 billion through 2032, through time spent on platforms, ad targeting and personalization.

## 6.7 Search, Recommendations Are Better Aligned

Using AI to improve search can boost gross merchandise value (GMV) and conversion rates as shoppers can more easily find items and use abstract keywords. Sifting through an endless aisle is tiresome, but offering the ability to narrow choices based on preferences and personal characteristics can change shopping behavior and boost sales.

Not only are retailers using large language models to filter across millions of products, but they're also able to request more details before providing search results. Lululemon's enhanced search functionality refines customer queries by prompting users for additional preferences, creating greater accuracy in meeting expectations.

Rent the Runway, Revolve and ThredUp are also using Gen AI to bolster search. Google's enhancements to browse with the technology -- circle to search and multisearch -- lets users remain in an app and search from within by circling a picture or video. In 1Q23, eBay said its search improvement was already boosting conversion and could eventually add about $1 billion in annualized incremental GMV.

BI **Cutting down the number of identical listings**

In 2Q, Etsy's gross merchandise sales for demi-fine jewelry grew 9% due to better search results and data-feed curation, and as the company continues to improve search and reduces the dominance of any one seller's item on the first page, it can continue to bolster conversion. Etsy has reduced the percentage of searches with two or more listings that may appear identical by over 70%, while the percentage of searches where a single seller's items dominated the page has been reduced by half.

Etsy's new AI-generated Gift Mode may attract shoppers and aid discovery for specific events and holidays, allowing consumers to enter information about who they're buying for and matching them with gifts on the platform.

Given Amazon.com's large scale (we break out its GMV in Fig. 40), it has a greater number of available data points to enhance search functionality, allowing for more personalized responses and product recommendations, further entrenching shoppers into its ecosystem. Through its large catalog of product attributes and customer-shopping information -- including preferences,

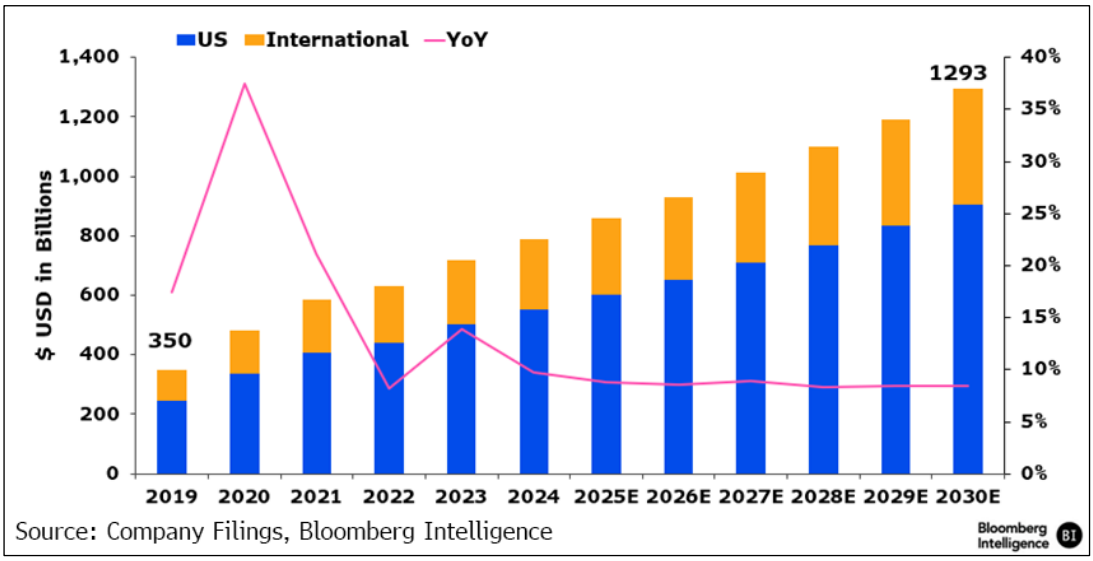

No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

search, browsing and purchase history -- the web giant can lean into large learning models to edit product listings to match customer preferences, making them more relevant for shoppers.

For example, if a customer is gluten-free, AI can more prominently display the relevant attributes of a product, which may have otherwise been listed at the bottom of a description.

**Figure 40: Amazon US, International GMV**

BI **Amazon racing toward $1.6 trillion in gross merchandise value**




Integrating recommendations and similar-item features can also boost conversion as it helps shoppers more easily find what they're looking for. Amazon's new AI personal shopping assistant, Rufus, can answer shopper's questions, compare products, help consumers find things for a specific occasion and discover information about a specific product without the buyer needing to read all the details. This, coupled with its customer review summaries feature, can help people find items faster.

## 6.8 Supercharging the Entire Supply Chain

Generative AI offers a big opportunity for retailers to improve supply-chain efficiency and costs, with Amazon.com already quickening inventory processing by as much as 75% and Nordstrom citing a 20% increase in distribution-center productivity. Demand forecasting, warehouse management, transportation and logistics and supplier-relationship management may all benefit, and omnichannel fulfillment is likely a key area for enhancement.

Amazon continues to deploy generative AI across its ecosystem, with fulfillment centers using more robotics systems to improve speed and efficiency. It has partnered with Covariant to build advanced AI models for its warehouse robots and deployed its robotic system, Sequoia, to help identify and store inventory at its fulfillment centers faster. Packages then move through an AI-powered robotic arm trio-- named Robin, Cardinal and Sparrow -- to sort, stack and consolidate items.

Amazon recently launched a Vision Assisted Package Retrieval tool for its delivery vans to find the correct packages faster, shortening routes by 30 minutes.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

Nordstrom's chief supply-chain officer said automated fulfillment centers are 5x more efficient than those without the tech upgrades. These enhancements also help reduce out-of-stock items.

The RealReal and eBay are also using AI to streamline the apparel and accessories authentication process, while Warby Parker and Etsy use generative AI to increase the productivity of their software engineers.

As e-commerce sales climb, returns are also rising, which may create bottlenecks in retailers' supply chains, notably when there's no dedicated return center. It often results in a buildup of inventory at warehouses, and as that merchandise waits to be processed, it depletes the value and possibility for a resale at full price. Retailers are leveraging AI to help cut the amount of returned goods, such as H&M analyzing which items get returned the most to help optimize for products that shoppers are more likely to keep, or Amazon's "predict-the-fit" feature to help shoppers select the right size, making them less likely to return it.

Returns are a growing problem for US retailers, with $890 billion of merchandise likely returned in 2024, or 16.9% of annual sales, based on the National Retail Federation.

## 6.9 Time-to-Market Speed Increase for Product Development

Growing use of generative AI to help with product design can help shorten the window from conception to production, reducing long lead times that often produce inventory imbalances and missteps. Product capsules can then be created much faster. Revolve, Nike and Lululemon, among others, are using generative AI to improve and speed up the product pipeline.

In addition to using AI to design marketing ads, notably billboards, Revolve is also using AI to produce new product lines. Last year it worked with winners from the inaugural AI fashion week to create a limited-edition, AI-created capsule collection to be sold on its platform. Though the implementation of generative AI is still in early stages, its growing use can continue to help designers broaden the possibilities more quickly for the future.

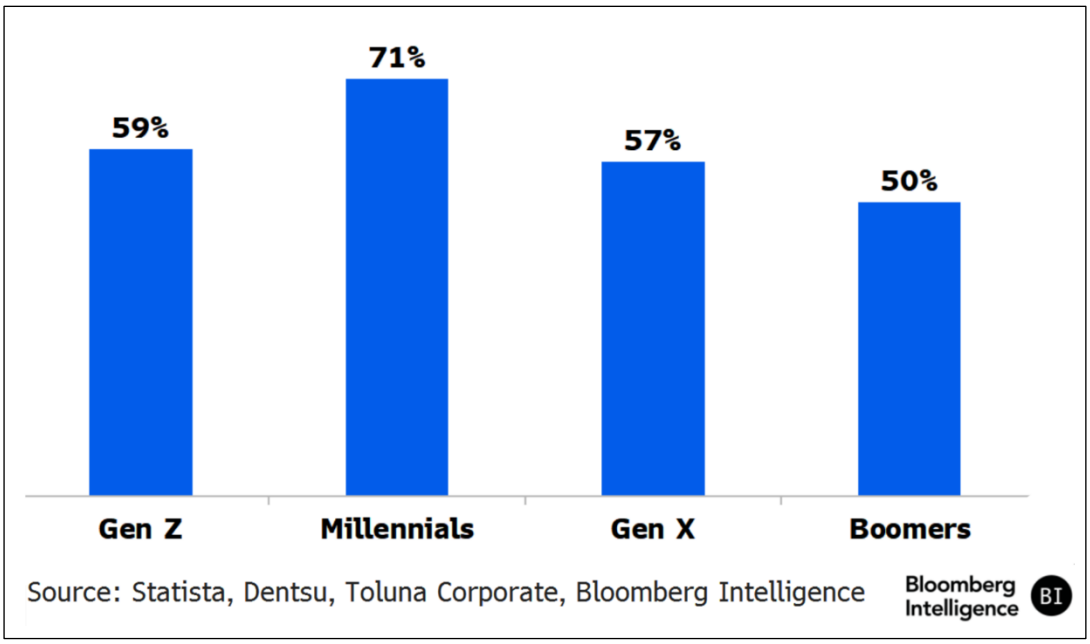
The use of generative AI tools, virtual reality and 3D rendering software has enabled Nike designers to create unique designs that reflect the personality of the athletes, adding newness to the product pipeline. The use of AI isn't yet a means to replace the designer but rather helps fuel their creativity with results available in minutes and seconds vs. weeks and months. The use of AI models to extrapolate data from texts, images, video and code, to create personalized products that fit the needs and molds of unique feet, can help Nike add more differentiation to its product lineup.

Hyper-personalized shoes may not make their way into mainstream retail just yet, but they can solve athletes' needs and reflect their persona better, helping create a halo effect for a version of the style when sold.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 41: Consumer Support for Brands Using AI-Aided Design**

BI **Younger consumers more supportive of brands that use AI-aided designs**




## 6.10 Improving Customer Service and Experience

Retailers (as shown in Fig. 42) are using AI to improve customer-service offerings by leveraging chatbots that can cut expenses and improve the buyer experience. Amazon's AI shopping assistants can answer personalized questions, while Wayfair uses AI to match shoppers to the right customer service agent.

According to IBM, using chatbots can reduce expenses 30% and boost customer satisfaction by 25%. Besides Amazon's AI shopping assistant, other companies employ chatbots to provide shoppers with immediate help. Etsy is also using the technology to streamline processes, including in trust and safety, for its engineers to drive efficiencies.

Wayfair uses gen-AI to match shoppers to the best customer service agent, predict what kind of service is needed by a customer with a complaint and what kind of discount might persuade a buyer to keep an item. Warby Parker is using gen-AI to quickly transcribe eye-wear prescriptions and help conduct virtual vision tests.

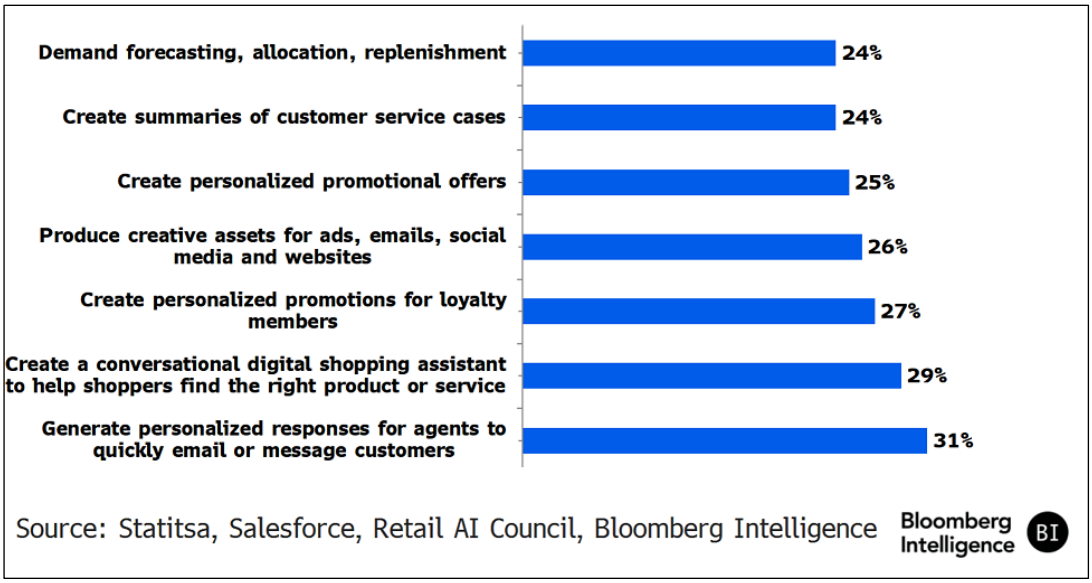
Amazon’s gen-AI-powered shopping assistant, Rufus, can answer customers' questions to help search for products, product comparisons and reviews to make recommendations to improve conversion. The company is also using AI to summarize customer reviews, which can spur conversion as it saves time for shoppers.

Amazon’s AI-enabled smart-home assistant, Alexa, has been programmed to understand complex natural-language queries and offer personalized recommendations, provide summaries of information and answer open-ended questions like “what to wear to the movies?"



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 42: Main Use Cases of Generative AI in Retail Industry**





No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Section 7. Hyperscaler Capex

## Hyperscalers Shift $258 Billion to Inferencing

DeepSeek's AI breakthrough will precipitate 10% annual growth in hyperscalers' capital spending total to $525 billion through 2032, our scenario analysis shows, as companies like Apple, Meta and Microsoft shift spending to focus on inferencing vs. pre-training several years earlier than forecast. Prior to DeepSeek, capex was projected to expand mid-single digits from $258 billion in 2024. AI-training outlays, which had been 40-50% of capital expenditures, may shrink as most foundational-model companies pivot to reasoning models at inference time. That could make inference-time cloud sales the fastest growing segment in the generative AI market, which we expect will grow to $1.8 trillion by 2032, as Project Stargate and other sovereign funds ramp up spending.

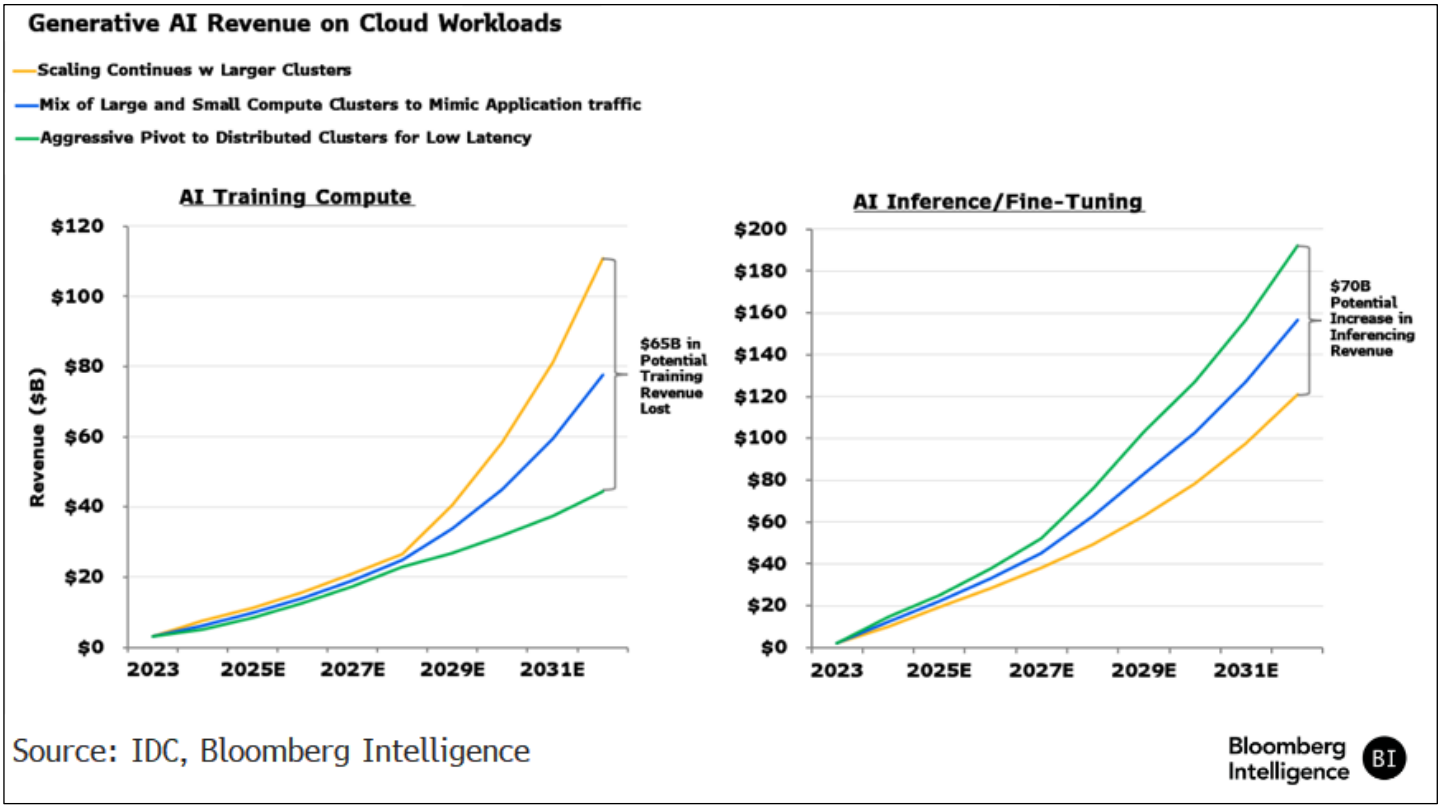
## 7.1 Rationalized Training Capex Likely

Capital spending growth for AI training could be much slower than our prior expectations (see Fig. 43), as the DeepSeek R1 release is likely to result in foundational-model companies moving to algorithmic improvements and increased investment in inference-time compute. The preference for reasoning models -- like DeepSeek R1 and OpenAI's o1 and o3 -- bodes well for inference-time scaling models compared with training.

Growth for training compute on the cloud may moderate to about a mid-to-high teens percentage, compared with a compound annual growth rate of around 50% in our prior forecast. In contrast, inference-compute growth could accelerate as large-language-model providers pivot toward reasoning and AI agent deployment. With pressure likely to increase on hyperscalers to deliver higher a return on investment for their capital expenditures, inference cloud revenue will likely take a higher priority over more upfront spending on pre-training. Our cohort of big-tech companies includes Microsoft, Amazon Web Services, Google, Meta, Oracle, Apple, Alibaba and Tencent.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 43: Bloomberg Generative AI Forecast**




BI

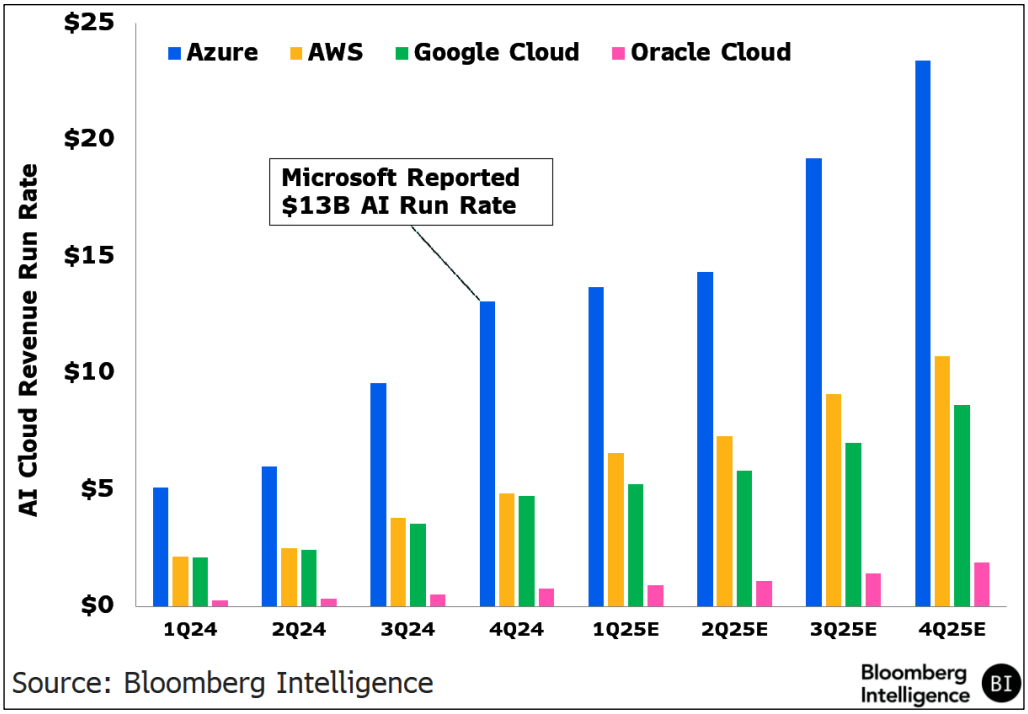
**Inferencing spending may rise, training fall**

Growing traction for OpenAI's o1 and o3, along with the open sourcing of DeepSeek R1 in January, is likely to push companies to allocate more compute for inference-time reasoning than on brute-force pre-training of large-language models. In addition, the Mixture-of-Experts approach used by DeepSeek R1 is likely to be incorporated by other foundational models to drive cost efficiencies.

With Microsoft's AI inferencing revenue already at a run rate of $13 billion, we expect inference-time cloud workload revenue could become the fastest-growing segment in BI's Gen AI forecast of $1.8 trillion by 2032.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 44: Hyperscale Cloud AI Revenue Run-Rates ($ Billion)**




BI

**Revenue poised to accelerate**

Though most foundational model efforts, including Microsoft-OpenAI ChatGPT, Google Gemini, Meta Llama, Anthropic Claude and xAI Grok, had plans to build bigger compute clusters for pre-training, we believe that could take a back seat as compute efficiency becomes a focus. Optimization techniques used by DeepSeek R1 to distill larger models could pave the way for lower spending on pre-training.

Some of the GPU compute on the cloud will be used for inferencing workloads to better match revenue generated with capital spending.

## 7.2 Hyperscaler Capex Growth Cliff Risk Tied to Training

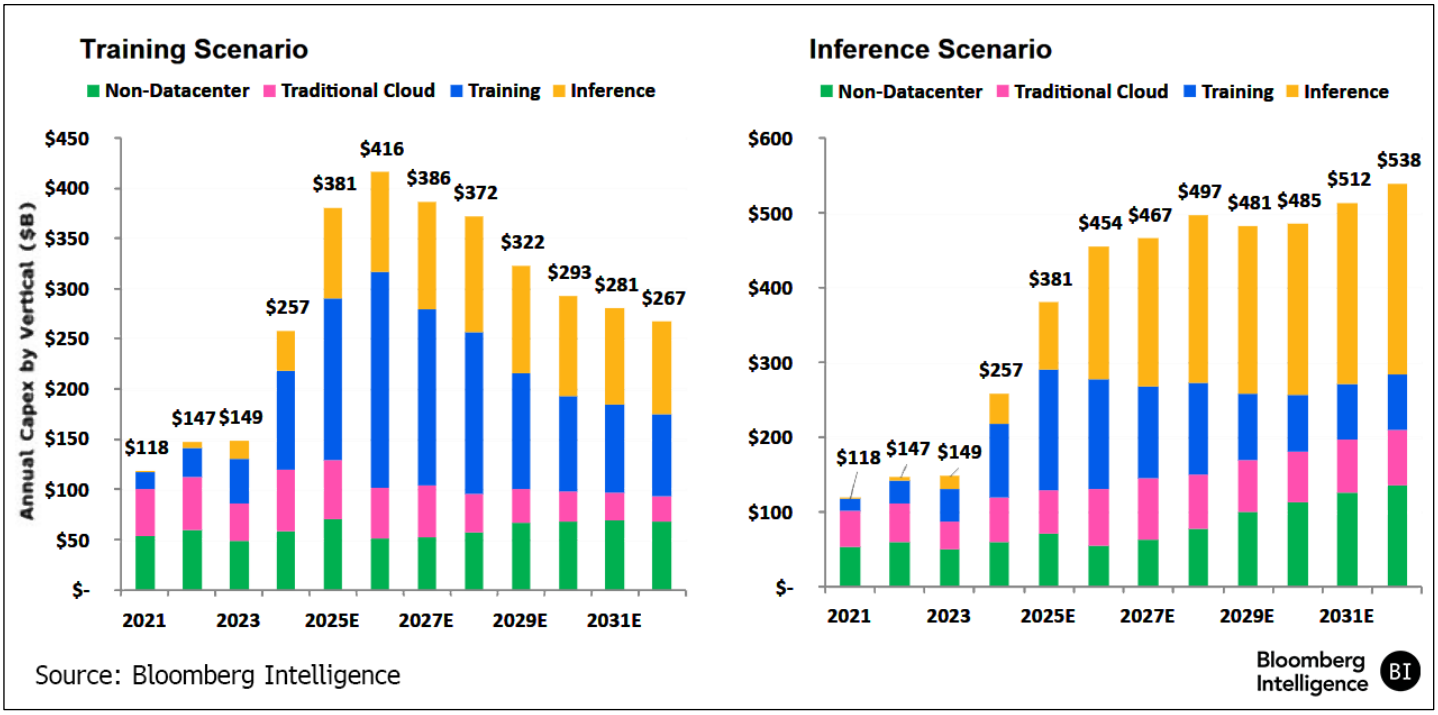
The risk of a capital-spending growth cliff is higher with continuous outlays for training clusters than in a scenario in which hyperscalers shift early to inference workloads. Alphabet -- with the use of its own tensor processing units (TPUs) for both training and inferencing -- may have an advantage to shift its capital expenditures quickly while others, including Microsoft-OpenAI and Meta, are more reliant on Nvidia for training.

Most hyperscalers may seek to optimize inference capital spending for latency and cost efficiency compared with the pre-training scaling focus, which requires bigger compute clusters for continuous data processing.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 45: Inference Scenario ($ Billion)**

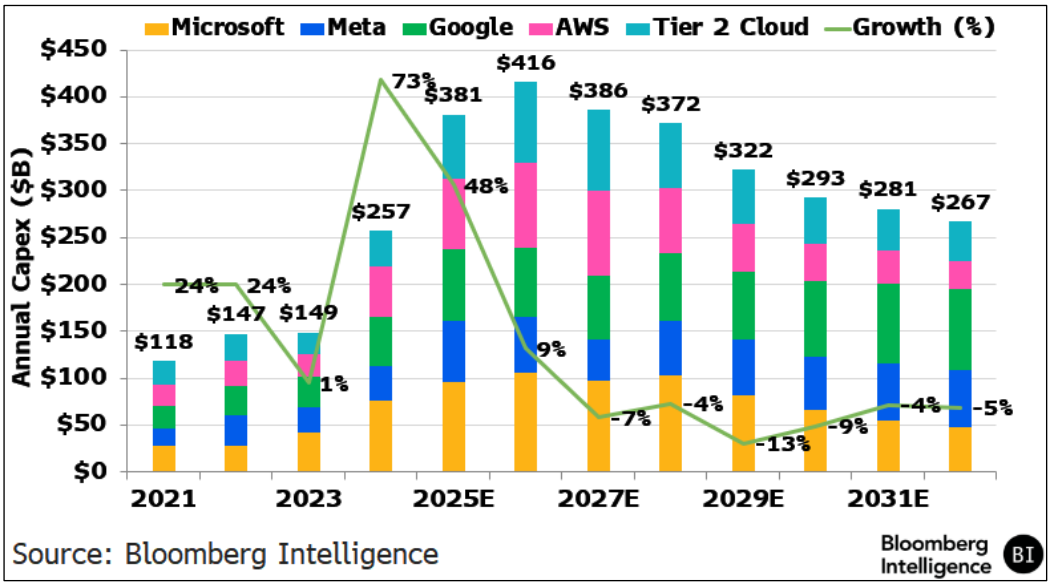



There's a higher chance of such a capital-spending growth cliff around 2027-28, we believe, if foundational model companies keep investing in larger compute clusters to train the next version of their large language models. xAI has the largest compute cluster with over 100,000 GPUs, closely followed by Meta, Microsoft-OpenAI, Alphabet and Amazon-Anthropic, all of which are investing in scaling the pre-training of their models. Training makes up about 40-60% of data-center capital spending for hyperscalers, which we expect to taper to 20-30% through 2032 to align the expenditures with revenue generation (Fig. 46).

BI

**Spending cliff possible in 2027-28**

**Figure 46: Training Scenario**



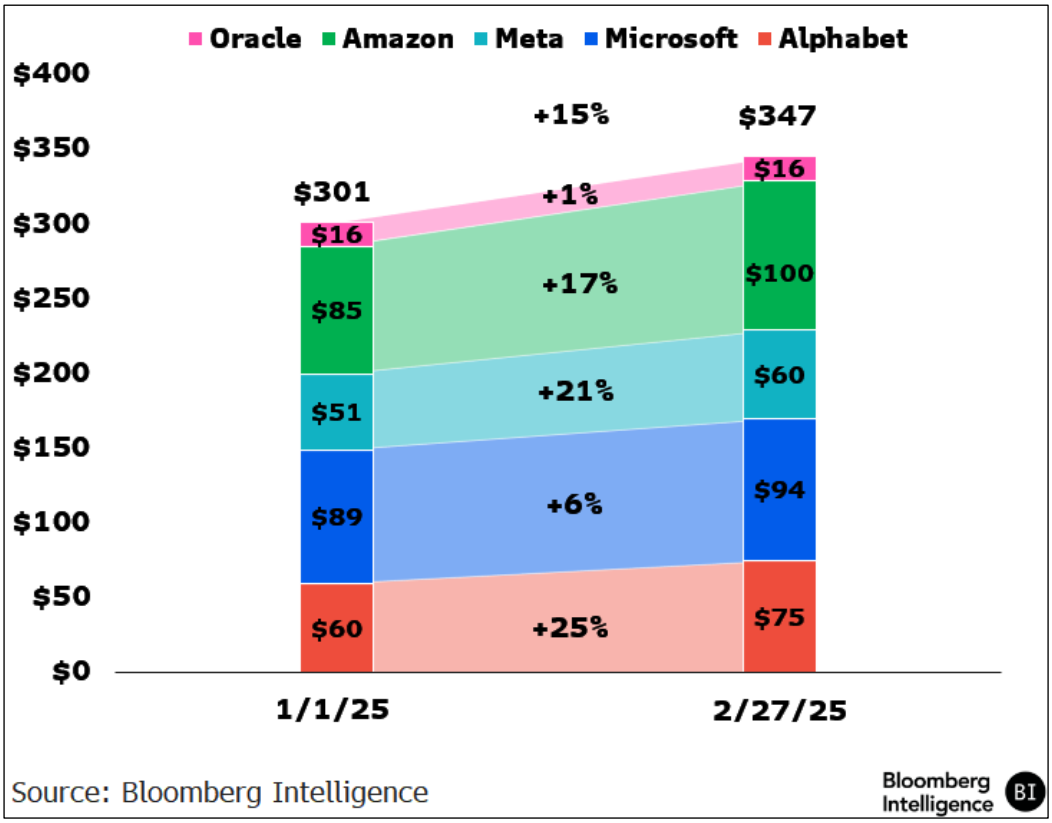



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

Cost efficiency and lower latency during inferencing could help determine the pace of capital-spending growth as companies look to match outlays with revenue generation. Data-center production for those focused on inferencing is expected to continue through 2032, given its cost has come down by about 90%. Total spending from hyperscalers could be around $80-$100 billion lower in a scenario where training remains a major part of data-center capital outlays.

**Figure 47: 2025 Capex Expectations Rose 15% Year-to-Date ($ Billion)**

BI **Sudden surge in capital spending expectations**




Among peers, Alphabet relies least on Nvidia GPUs for training its Gemini model, and it hasn't invested in building a larger cluster of compute in one data center, like xAI, Meta and OpenAI. Amazon.com has also been investing in its Trainium and Inferentia chips to focus on cost efficiencies for its data-center operations. The utilization of data centers with large compute clusters could threaten the hyperscalers' ability to sustain capital spending growth beyond 2026.

Hyperscalers with their own foundational models spend about 60% of capex on training, while inferencing spending is split between their own applications and enterprise consumption through their cloud segments.

## 7.3 Microsoft to Lead as Big-Tech Capex Heads to $381 Billion

Recent comments by top tech companies about higher AI spending suggest that total capital expenditures in 2025 could reach $371 billion, a 44% increase over last year (see Fig. 48).



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

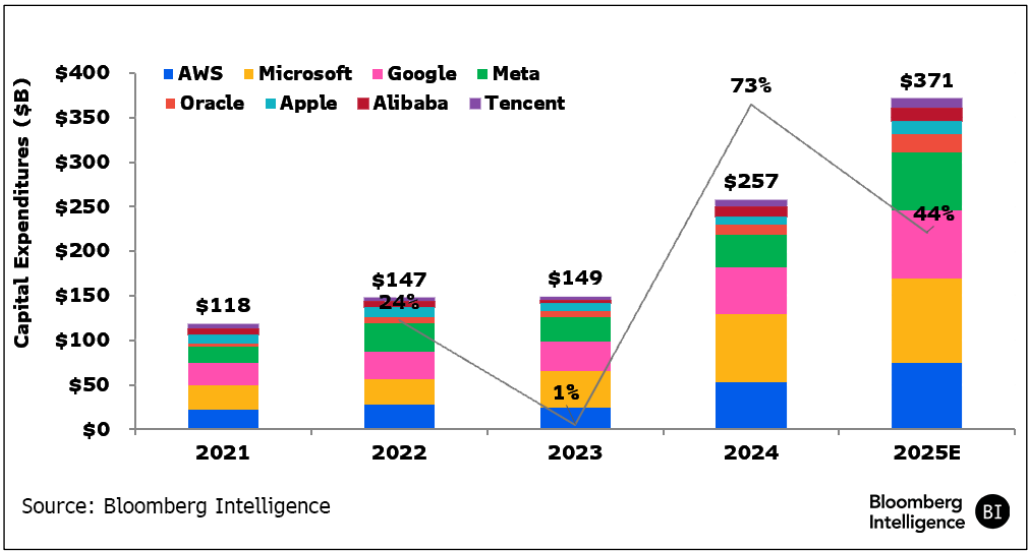
Microsoft CEO Satya Nadella's talk of leasing more than building in 2027 and 2028 tempers optimism that capital spending growth rates will begin to slow beginning in 2026.

Meta CEO Mark Zuckerberg's announcement of $60-$65 billion in spending was followed by promises of higher outlays by Google and Amazon.com to ensure enough capacity to power their AI ambitions. All these communications came out after SoftBank, OpenAI and Oracle announced a joint venture to lay out $100 billion in 2025 and $500 billion over a four-year period for the Stargate Project.

Our analysis doesn't include Stargate outlays, which may change once the joint venture's details are formed. This figure could be up to $346 billion in 2025, excluding Chinese participants.

**Figure 48: Bloomberg Tech Capex Analysis**

BI **Big tech companies all in on AI capex**




Expansionary capital spending could extend to around 75% of total capex by 2026, according to our analysis, with the bulk of this going toward GPUs, CPUs and servers, followed by storage, networking, cooling and utilities. This spending will likely take place in the next two to three years to meet the high demand for AI services, with maintenance starting to take back its share of overall capex after. We're less concerned about these companies overspending in the near term, because this hardware can eventually be repurposed for inferencing or more general cloud workloads.

Microsoft stood out as the lone hyperscaler to stick to its capital-spending goal ($80 billion) for AI data centers in fiscal 2025 (ending in June 2025), with the total likely topping $87 billion, or 57% over 2024. Extrapolating the spending pace to year-end, we calculate total outlays of $96 billion for calendar 2025.

Microsoft allowing OpenAI to seek its model training with other cloud providers is another example of an action taken to potentially lower spending needs.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

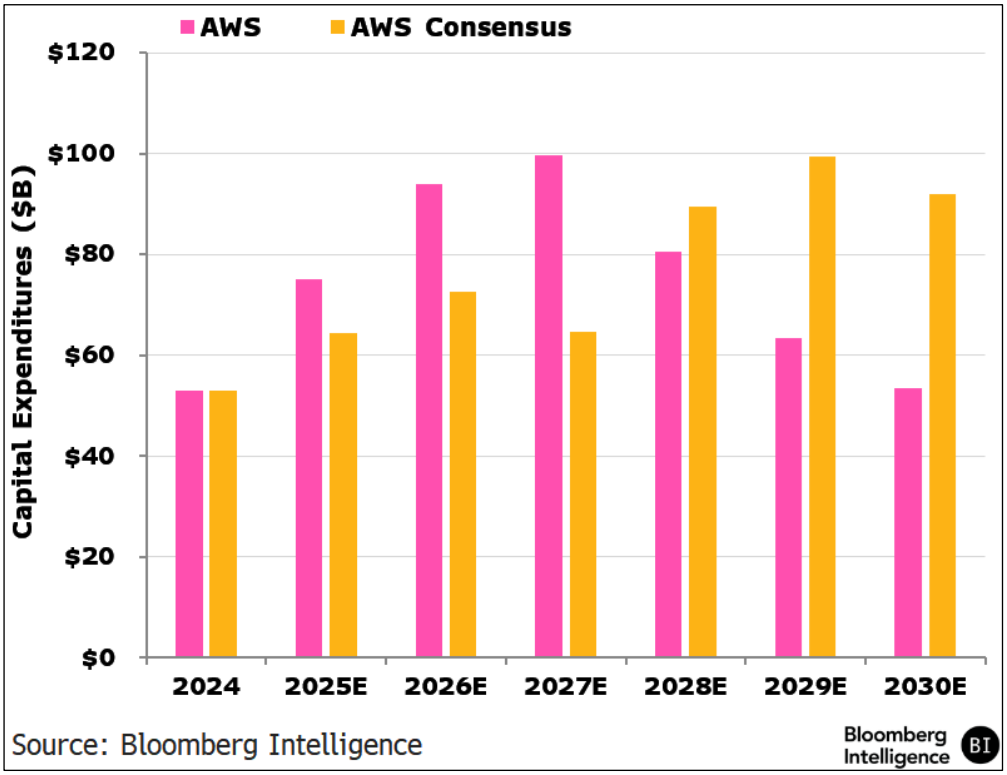
## 7.4 AWS, Apple Capital Spending Also Going Up

Amazon.com said on its fourth-quarter earnings conference call that it plans to maintain the same pace of quarterly capital spending in 2025, which we calculate to be around $105 billion. Of this, we anticipate roughly $75 billion will be allotted to Amazon Web Services (AWS), about $11 billion more than consensus (Fig. 49). Given AWS' high market share in cloud computing, at over 40% of the total, we don't expect management to back down from even higher outlays if needed, yet we also don't anticipate a noticeable change in sales growth in the near term because AWS lacks any major consumer AI application running on its platform, unlike Microsoft, which powers ChatGPT's back-end for inferencing.

**Figure 49: AWS Capex BI Analysis vs. Consensus**

BI

**Consensus for AWS capex may be high**




Apple's capital spending appears poised for a significant increase -- at least over the next two years -- with an additional $500 billion investment announced in the US for the following four years. Our calculations point to around $15 billion in 2025, up from $10 billion in the prior year. We see the lion's share going toward data centers capable of supporting more intensive Apple Intelligence inquiries, which over time could reduce the company's reliance on third-party AI model providers like OpenAI for questions that require greater reasoning.

This capex spike could come at the expense of its share buybacks, which distributed close to $100 billion in the past 12 months.



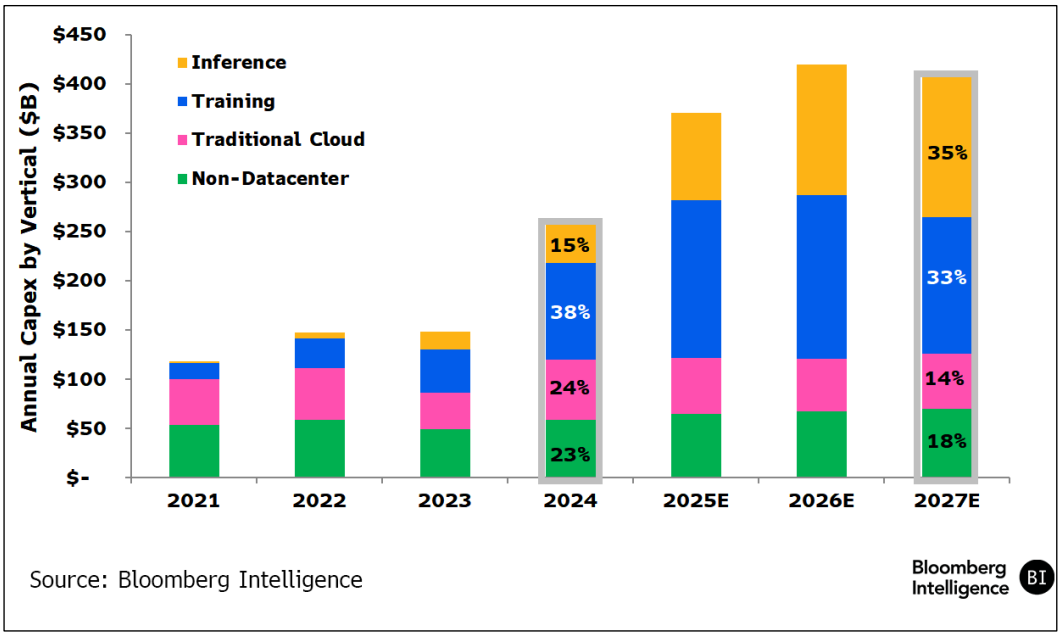
No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

## 7.5 Inferencing May Surpass Training Earlier

Compute consumption for inferencing is likely to surpass training workloads at least three years sooner than our initial projections for the generative-AI market (see Fig. 50). That view had training as the bigger aspect than inferencing, given the scale requirements for the former. The biggest driver for the change in our assumptions: The pressure on hyperscalers to generate return on assets for their capital spending and the pivot toward focusing on cost efficiencies for inferencing after DeepSeek's release.

Though most hyperscale cloud vendors are building big clusters of size (100,000-200,000 GPUs) in data centers for training the latest versions of their large language models, we believe the efficiency of DeepSeek and low utilization of GPU clusters of that size could limit further scaling for training LLMs. In contrast, inferencing requirements may expand as a result of the higher efficiency exhibited in the DeepSeek R1 release.

**Figure 50: Inference to Outspend Training by 2027**




BI **Capital outlays could be at near-term peak**

Most hyperscalers use the latest GPU chips for training the most recent version of their large language models, and cluster size helps drive a higher number of synchronous floating-point operations. The utilization of a large GPU cluster could be lower than the smaller clusters used for inferencing, where handling user-traffic fluctuations and low latency are the most important factors.

The use of custom ASIC chips could proliferate in the wake of DeepSeek's open-source release of its R1 reasoning model. Google is furthest along with its TPU chips, followed by Amazon Web Services with its Trainium chip to optimize workloads for cost and power efficiency. Broadcom's view of $60-$90 billion in serviceable addressable market by 2027, around three to five times higher than 2024, suggests that hyperscalers' focus on scaling inference will drive continued investment in custom silicon and networking.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

Pre-training of large language models involves high upfront investments, and we believe the breakthrough in reasoning models from OpenAI and DeepSeek could lead to higher monetization for hyperscale providers. Microsoft's Azure segment, used for running AI inferencing workloads, already has a 12-to-14 percentage-point revenue contribution from generative AI, while Amazon Web Services and Google Cloud are likely logging a 5-to-7 percentage-point addition, based on our calculations.

## 7.6 Stargate, Sovereigns Spending Key to Optimism

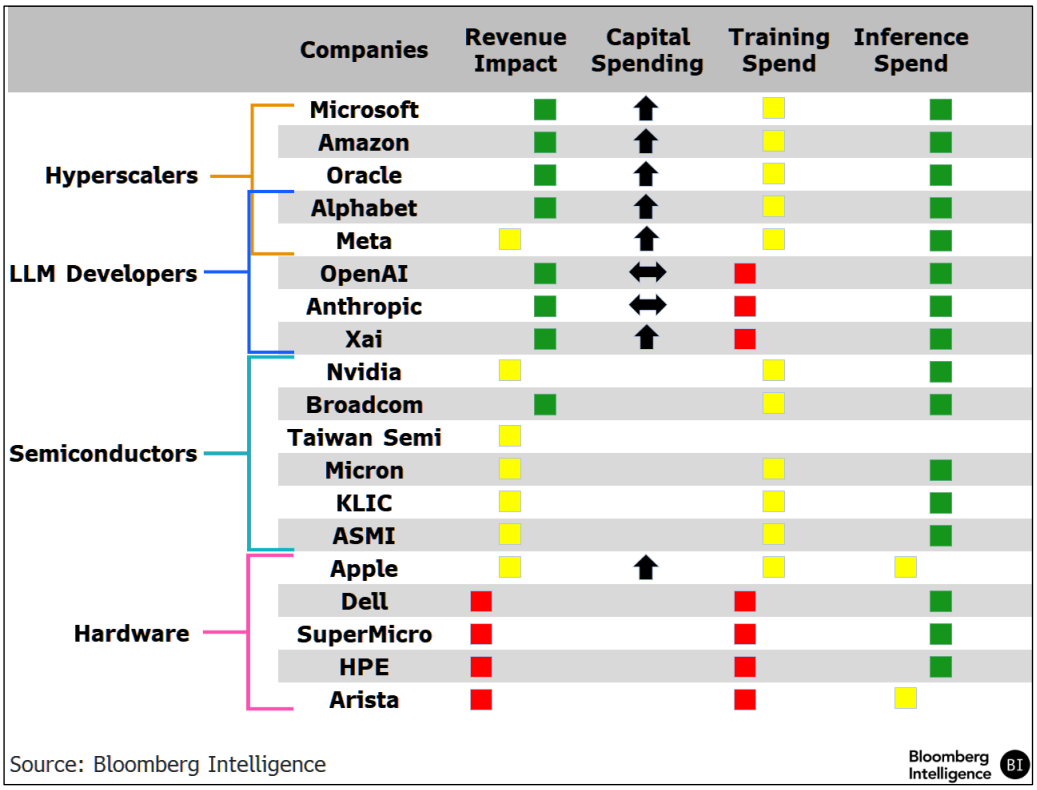
Increased 2025 capital spending among hyperscale-cloud providers to about 40-50% growth, along with the additional data-center build-out for Project Stargate and European Union funding, suggest a longer tail for the cycle. Though a shift toward smaller distributed clusters could be a headwind for capex growth after 2025, sovereign funds could offset that with incremental spending.

Prior to the DeepSeek R1 model release, most companies had road maps for building a 1-million-compute cluster to train their foundational models that would raise capital spending for hyperscale cloud providers. But an increased preference for smaller distributed clusters after DeepSeek's breakthrough could now limit the size of training clusters.

BI

**DeepSeek has catalyzed shift from training to inferencing**

**Figure 51: DeepSeek Impact Table**




Though sovereign wealth funds were about a mid-single-digit contributor to incremental data-center expansion in 2024, Project Stargate in the US and the EU's proposal for a €200-billion fund suggest a long tail for the capital-spending cycle. Sovereign spending makes up about low



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

teens of Nvidia's data-center revenue run rate of $125 billion, which could get a significant lift from the Stargate and EU spending for building data centers.

Capital spending from Stargate and the EU sovereign investments could add about $100 billion to the $371 billion in expected aggregate hyperscale capex for 2025, we calculate.

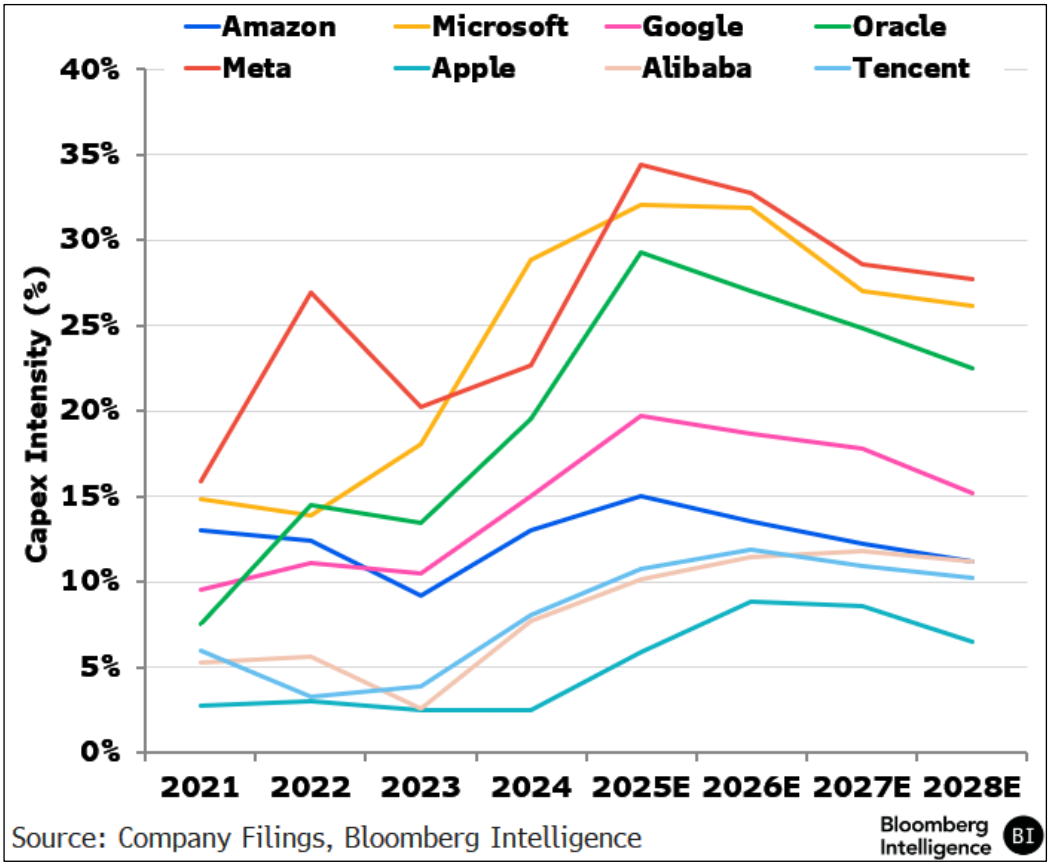
Most hyperscalers have seen an increase in capital spending intensity to expand data-center infrastructure for generative-AI compute for both internal consumption and deploying customer workloads on their public clouds. While capex intensity for Meta has risen to over 30%, Microsoft and Google are at about 20%, their highest in the past 10 years.

The pressure for return on investment is expected to increase, with revenue generated less than 10% of the total aggregate capex of around $258 billion for gen AI in 2024.

**Figure 52: Hyperscale Capex Intensity (%)**

BI

**AI inferencing demand drive capital-spending intensity higher**




## 7.7 Traditional Cloud Is Likely to Get Boost from AI Inferencing

Hyperscale cloud is likely to quickly pivot toward AI inferencing in favor of training workloads, which we believe will help boost segment growth. Though Microsoft's AI contribution to its cloud revenue is the highest at about midteens, a $13 billion run rate, we expect other hyperscalers to contribute revenue of over 10% by 2025, primarily driven by AI inferencing.

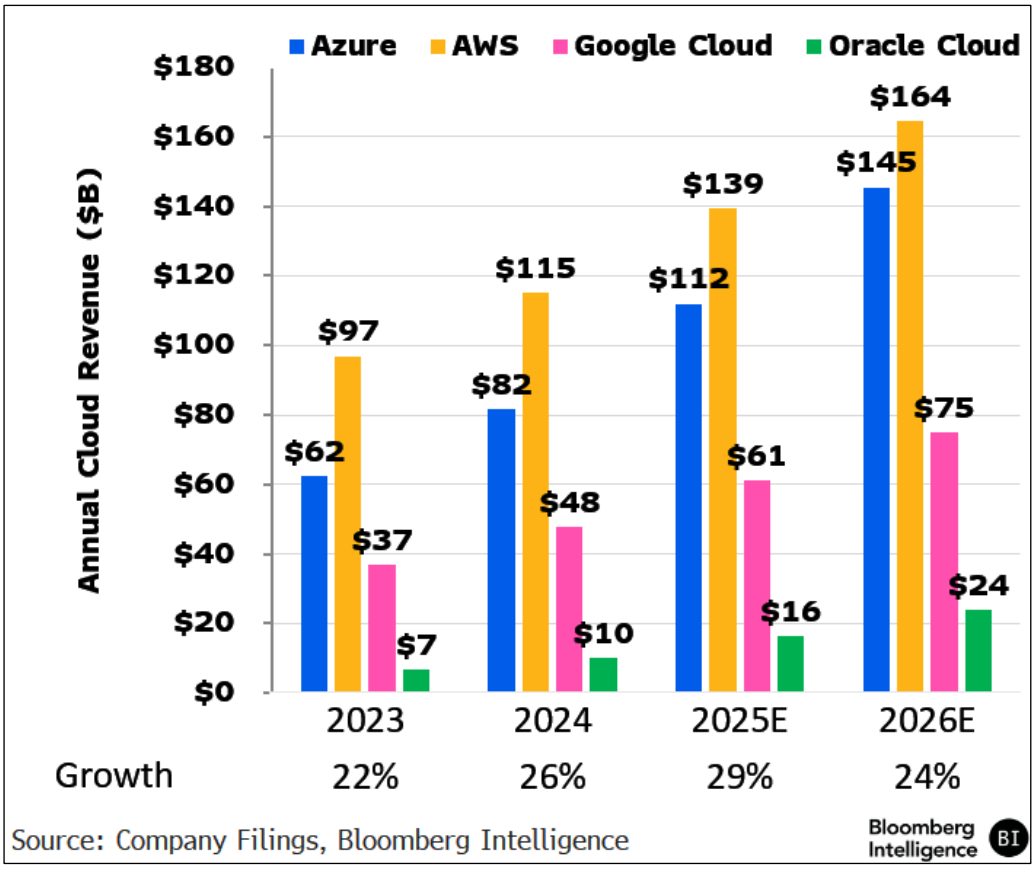

No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

Microsoft's high contribution from AI comes from traction with deploying OpenAI's model on its cloud. Cost per million tokens fell by about 90% in the past 12 months across all the major foundational model providers.

AI inferencing workload revenue could rise about 10x from about $20 billion now, making it one of the largest segments in the generative-AI market.

A boost in AI-inferencing workloads could further aid consumption of traditional cloud features including databases, cybersecurity and observability (see Fig. 53). We don't see any substitution of traditional cloud workloads with AI -- the use of a chatbot or natural language interface is likely to drive more engagement with applications and potentially boost daily active users.

**Figure 53: Cloud Growth Above 20%**




BI

**Cloud demand could get a coincident lift**

With the shift to inferencing, revenue growth in hyperscaler cloud segments could move in line with capital spending increases, as reasoning models with inference-time compute could charge enterprise customers for consumption. Most hyperscalers might look to expand their GPU inferencing capacity across distributed data centers to improve latency instead of looking to build the largest GPU cluster for pre-training.

So far, hyperscale cloud vendors haven't seen much pressure on margin, even as they expect AI chips to depreciate faster than traditional CPUs, with 2025 estimates falling one to two percentage points for AWS, Meta and Microsoft Azure, holding steady at Oracle and rising six



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

percentage points for Google Cloud. The shift away from pre-training to inference scaling could lower potential margin headwinds from higher infrastructure and power-related costs.

Though scaling laws have focused on pre-training so far, we believe the DeepSeek model release and the higher accuracy of reasoning models, such as OpenAI's o1 and o3, have resulted in a shift toward inference-time compute. Based on the average capital-spending increase of about 55% for the hyperscalers, per consensus and recent capex views, we expect companies to channel a large portion of the data-center spending to ramp up AI inference capacity.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Section 8. Processing, Memory Chip Demand

## Growth Opportunities Spread Around the Globe

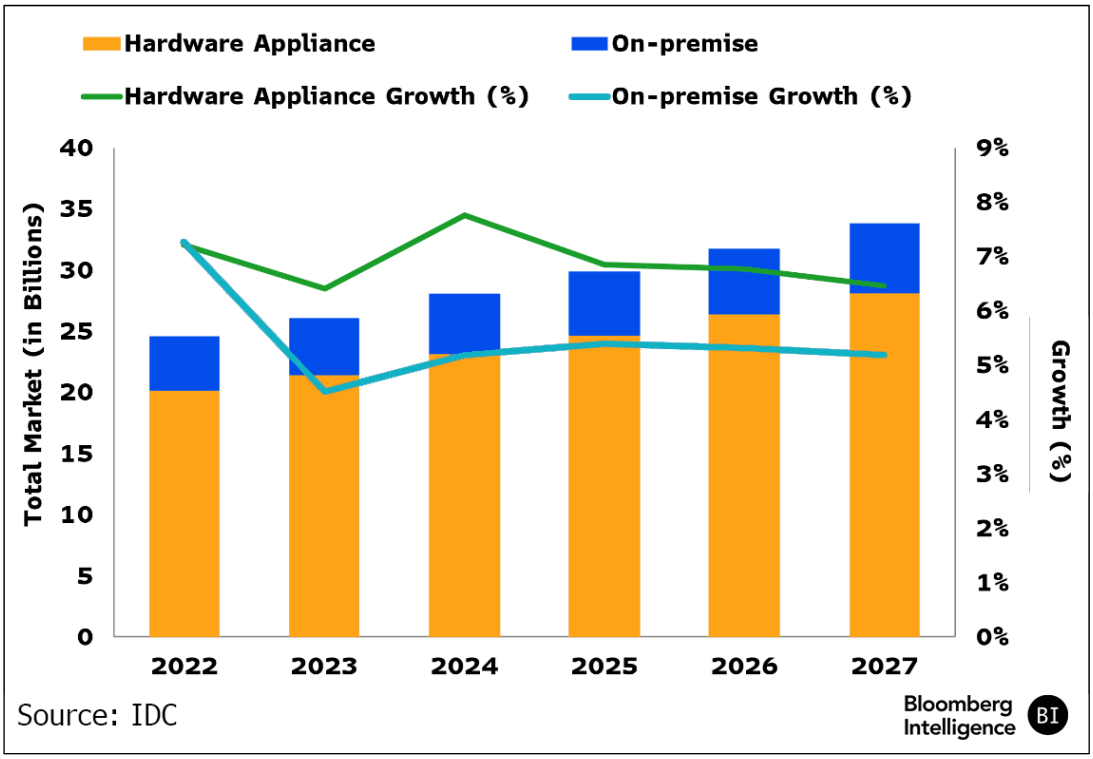
The spread of artificial intelligence could drive demand for graphics processing units (GPUs) and dynamic random-access memory (DRAM), lifting sales at SK Hynix, Samsung Electronics and Micron. We believe memory semiconductors will play a key role in the expansion of the data-center chip market, along with AI accelerators, with each expected to grow more than 15% annually over the next three to five years.

## 8.1 TSMC Prowess Sets It Apart From Pack of Rivals

Fortinet and Palo Alto Networks may leverage custom semiconductors for generative AI, supporting steady refreshes of their hardware and software firewalls. Fortinet added software-defined wide area network (SD-WAN) functionality to its ASIC chips, helping it to take share from traditional firewall vendors such as Cisco and Check Point. Palo Alto successfully bundled its Prisma, Cortex and virtual firewalls to help enterprise customers secure their on-premise and public-cloud workloads.

**Figure 54: Network Security Market**

BI **A consistent refresh of hardware and software**






No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

AMD's energy-efficient AI accelerators stand to gain amid rising scrutiny over power consumption by data center AI chips.

TSMC is on track to maintain strong revenue growth, averaging 22% compound annual growth during 2025-2027, driven primarily by sustained demand for AI accelerators in both training and inference applications. Even if chipmaking orders slow from AI GPU designers like Nvidia and AMD, increasing orders from custom-AI chip developers, such as Broadcom, will fill the gap swiftly. The growing complexity of inference workloads in data center and multi-modal scaling will sustain global AI accelerator growth and enhance demand for TSMC's N3 and N2 manufacturing technologies, as well as its advanced 2.5D packaging solution.

TSMC's sales related to AI chips for data centers are anticipated to double by the end of 2025 and sustain annual growth above 33% through 2029, in our scenario analysis.

**Figure 55: BI Scenario Analysis**

| | | | | | BI Scenario | | |
|---|---|---|---|---|---|---|---|
| | 2020 | 2021 | 2022 | 2023 | 2024E | 2025E | 2026E |
| (Revenue By Applications, $Mn) | | | | | | | |
| **Smartphone** | 21,917 | 24,887 | 29,835 | 26,146 | 31,013 | 35,849 | 39,434 |
| *Sales Mix %* | *48%* | *44%* | *39%* | *38%* | *37%* | *35%* | *34%* |
| **HPC** | 14,938 | 21,045 | 31,295 | 29,992 | 39,460 | 49,416 | 59,480 |
| *Sales Mix %* | *33%* | *37%* | *41%* | *43%* | *47%* | *49%* | *52%* |
| **Others** | 8,632 | 10,904 | 14,856 | 13,220 | 13,824 | 16,101 | 16,202 |
| **TOTAL Sales** | $ 45,487 | $ 56,835 | $ 75,986 | $ 69,358 | **$84,298** | **$101,367** | **$115,117** |
| *Growth %* | | *25%* | *34%* | *-9%* | *22%* | *20%* | *14%* |
| **Sales Consenus Estimation (as of 31/12/23)** | | | | | **84,085** | **98,783** | **110,357** |
| | | | | | *21%* | *17%* | *12%* |

*Key Assumptions:*

**1. Smartphone related revenue to grow at an 15% compound average rate from 2024 to 2026**
**2. High Performance Computing (HPC), which include AI accelerator, server processors, PC will grow at a 26% compound average rate from 2024 to 2026**
**3. Total sales contribution from smartphone and HPC chip businesses will be 84% in 2023, 2024 and 87% in 2026**
**4. TSMC will still secure over 90% of global AI accelerator production orders**

Source: Bloomberg Intelligence

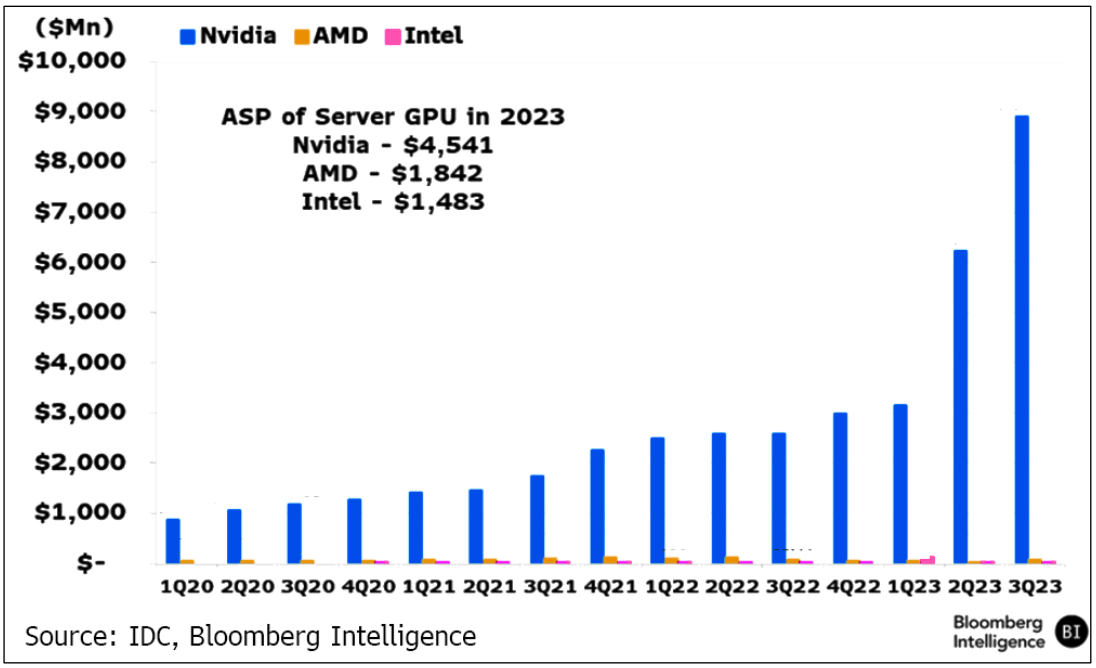
Bloomberg Intelligence BI

TSMC's dominance in leading-edge node semiconductor manufacturing processes positions it to maintain its hold on the lion's share of AI chip production orders from key players such as Nvidia and AMD. That advantage is expected to continue, thanks to the company's strong production yield. Also, many AI chip designers prefer TSMC's CoWoS packaging for its superior interconnection density, larger package sizes and cost effectiveness.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 56: AI Server ASPs**




BI

**Nvidia dominates GPU market**

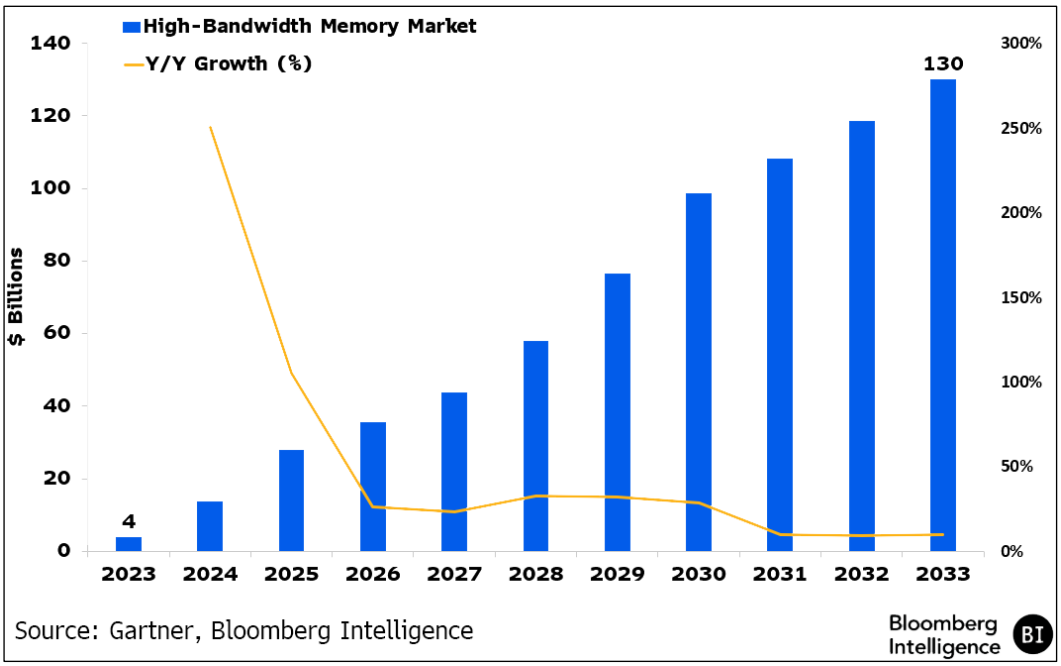
## 8.2 HBM Exposure to Inferencing Demand

High-bandwidth memory, an ultra-high-speed semiconductor that's critical to the development of artificial intelligence, can grow at a compound annual rate exceeding 40% to reach sales of $130 billion in 2033, from just $4 billion in 2023, based on our scenario analysis. (See **BI's HBM report**.) Part of the steep increase in sales will come from higher volume, but a large part will be due to new generations of HBM with greater functionality and higher prices. Since SK Hynix has built a good relationship with Nvidia on the current generation of GPUs, it can benefit from increasing demand HBM can be more easily customized than standard DRAM, so the company might protect its top market share by working closely on customized products with major customers such as Nvidia.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 57: High-Bandwidth Memory Market**




BI

**Servers the biggest user of DRAM; make up more than one-third of SK Hynix' total**

For Nvidia's new GPUs, Micron could also expand revenue, as it started volume production of HBM chips for Nvidia in early 2024. Results at Samsung, the world's largest DRAM maker, also may rise amid increasing use of GPUs and HBMs. The eventual switch to next-generation HBM4 or 4E chips could be an opportunity for Samsung to wrestle back market share. Customer certification of HBM4 may begin as early as 2025, and full-scale market penetration could occur in 2026.

Escalating use of AI for inference applications might heighten the value of DRAM formats like graphic double-data rate (GDDR), which is used in retail PC graphic boards where cost is critical, and low-power-consumption double-data rate (LPDDR), which is mainly deployed in smartphones.

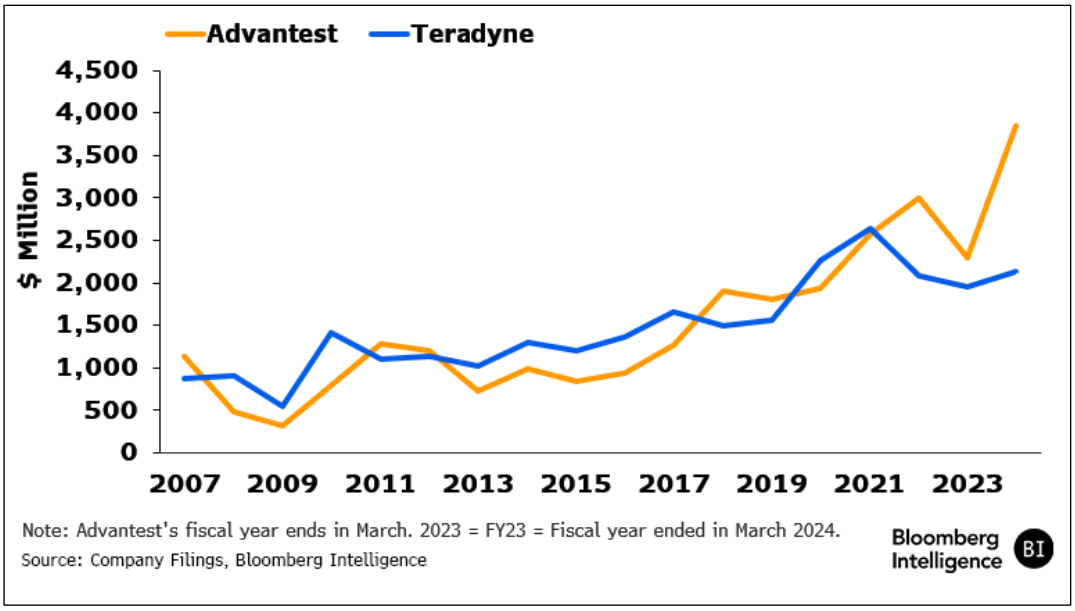
Power consumption and data telecommunications would become excessive if all processing were conducted on servers. That will make it necessary to perform AI tasks on edge devices, which could fuel a surge in DRAM orders for products including PCs, autos, robots, smartphones and security cameras, buoying sales for Samsung, SK Hynix and Micron.

The speed of chip performance improvement required for AI is faster than the evolution of miniaturization and advanced packaging, meaning quality isn't certain. As a result, the role of chip testers (see Fig. 58) that can accurately assess performance and quality might rise sharply. Teradyne has a strong competitive edge in the field. And customers have given high marks to Advantest's T5000 series of memory-chip testers and V93000 series of system-on-a-chip testers.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

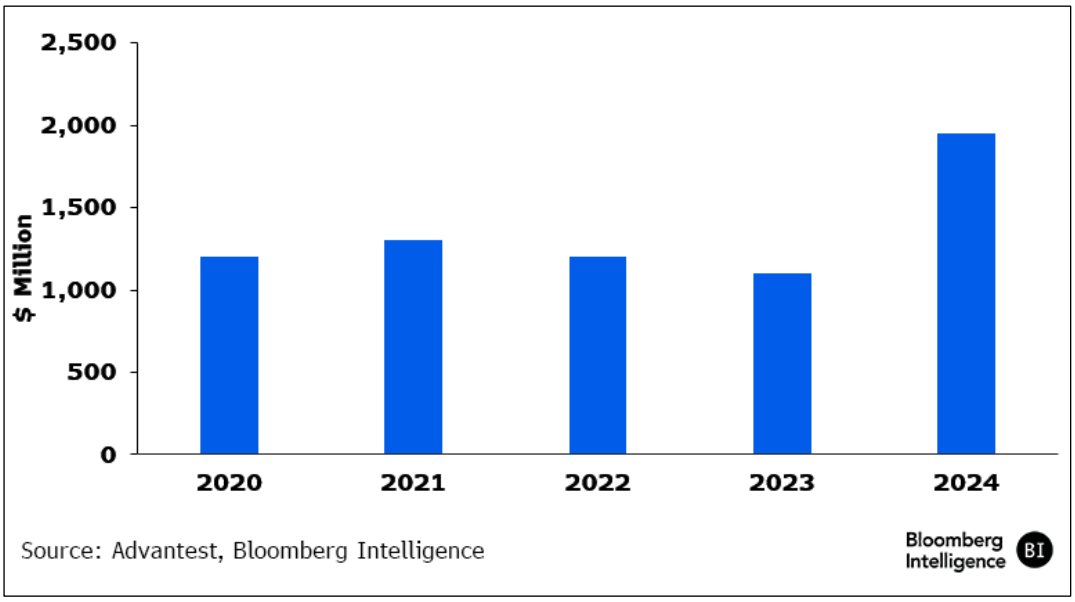
**Figure 58: Chip Tester Sales for Advantest, Teradyne**




**Quality assurance critical for single and packaged chips**

Quality assurance will be necessary not only for single chips, but also for chip packages, and entire system-level tests (SLTs) will become more important. Three-nanometer technology is taking off and chiplet packages, which place several chips or dies on a substrate vertically, could be adopted in a few years. Beyond identifying defective products, tester makers like Teradyne need to reduce the risk of misclassifying good products as deficient, requiring substantial technological capability. Advantest strengthened its SLT business by acquiring Astronics in 2019 and Essai in 2020.

**Figure 59: Memory Tester Market Size**






No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

## 8.3 User-Based Software Growth May Shrink

The onset of AI-powered agents and copilots could enable companies to optimize resources, build more applications in-house and that could reshape the seat-based businesses that sell to them. As AI-driven productivity gains broaden across industries and functions, job growth could slow, pressuring seat-based revenue models. We're also doubtful of a revenue uplift from price increases as AI features standardize.

Managed-services providers like Confluent that have built their business around open-source models may need to beef up their product innovation. The ease of code generation through simple natural-language commands -- which has become possible with the use of large foundational models and applications such OpenAI's ChatGPT -- potentially brings application development capabilities to the masses. It's similar to the content generation that became possible with the confluence of mobile phones, social-media platforms, broadband connectivity and easy content-editing tools. With code-generation workloads potentially easing, developers' time could be used to build more applications in-house, using open-source software and engines rather than through third parties or proprietary products.

**Figure 60: Shifting Business Models**

BI **Observability market shifts, and vendors need to follow**

| | | Companies Exposed To The Theme |
|---|---|---|
| Seat Based Models | AI could act as double-edged sword. While it does bring prospects for streamlining cost base, the productivity gains that come with AI and the standardization of AI features across technology platforms could potentially narrow product edge (differentiation) and curb the seat and headcount expansion led growth. | Atlassian<br>GitLab<br>Monday<br>Asana<br>Smartsheet |
| Opensource, AI Disruption | Adoption rates on solutions like Grafana in observability, OpenTofu in infrastructure provisioning, Apace Icerberg, PostgreSQL, Apache Spark, Trino for data workloads, Apache Cassandra, Weaviate and Milvus for Vector search -- need monitoring and could potentially become more pervasive and limit incremental growth fro managed servce solutions. | Confluent<br>MongoDB |
| Point Solutions | The integration of AI capabilities into existing technology solution and as large technology vendors enter into adjacent verticals, they could stand to reap from customer moves to consolidate technology stack and vendors. | C3.ai<br>UiPath |
| AI Infrastructure | The rising capex towards AI infrastructure could also drive up demand for infrastructure provisioning and solutions that help monitor and optimize those costs. | Datadog<br>Dynatrace<br>Elastic<br>Cisco<br>HashiCorp<br>Flexera<br>IBM (Apptio)<br>ServiceNow |

Source: Bloomberg Intelligence

Bloomberg Intelligence BI

We see risk to revenue consensus for point-solutions software companies, seat-based revenue models and those that lack scale and face narrowing product edges like Asana, MongoDB, C3.ai, Confluent, Rubrik and UiPath. The changing competitive landscape in the observability sector, with open-source solutions like Grafana gaining scale, could force companies like Datadog, Dynatrace and Elastic to sharpen their execution and demands monitoring.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

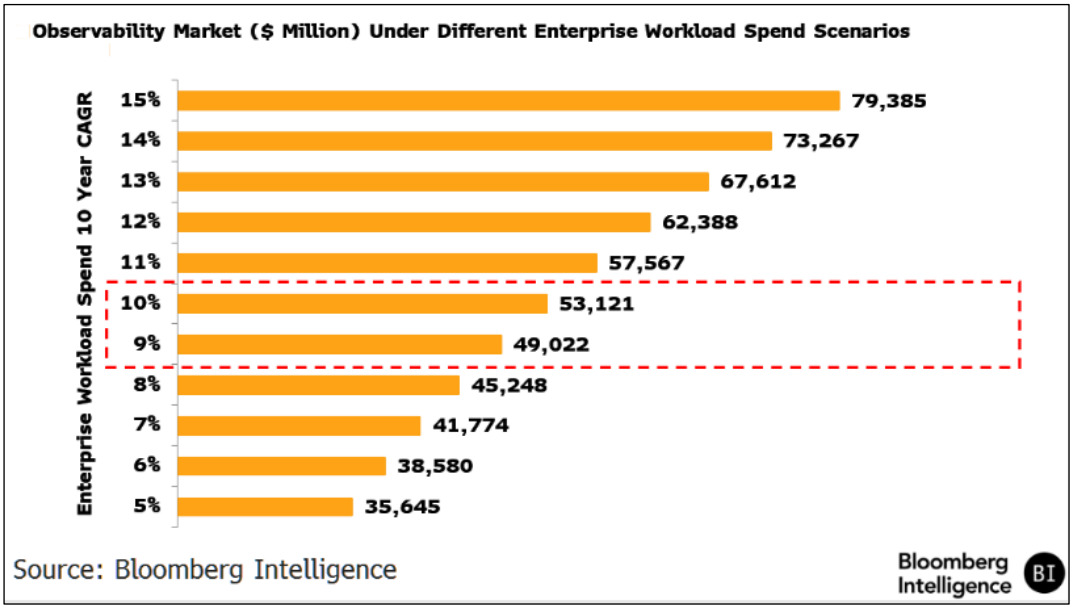
The criticality of digital assets -- application, compute, storage and data layers -- to modern business demands always-on monitoring to prevent operational disruptions and to remediate at speed. With AI's growing capabilities, the risk of bad actors could rise, driving greater demand for always-on monitoring solutions covering many layers of IT value chain.

Rising capital spending toward AI infrastructure could also drive up demand for infrastructure provisioning and products that help monitor and optimize those costs.

## 8.4 Observability on a Path Toward $50 Billion

AI-led applications, agents and underlying infrastructure growth could boost demand for monitoring, a fledgling vertical we see nearly tripling up to a $53 billion addressable market by 2033. LLM applications are resource intensive -- both in compute and memory -- and robust observability could aid efficiency and control costs. As AI applications advance from pilot to deployment and begin to scale, demand for solutions that monitor the performance of large language models can rise, while closer monitoring of LLM-fueled applications, which will likely power more real-time services, is critical to address service bottlenecks, avoid disruptions and root out possible AI hallucinations. Scaled-observability providers like Datadog, Dynatrace and New Relic have already launched solutions in this space. In July, Datadog announced the general availability of LLM Observability, which allows application developers and machine-learning engineers to monitor, troubleshoot and secure AI applications. This stands to benefit observability platforms geared for monitoring the health of enterprise applications and their underlying IT infrastructure, supporting business resilience and optimizing resources, and we see players like Datadog and Dynatrace best positioned to monetize, given their customer reach and product enhancements.

**Figure 61: 2033 Observability TAM Scenarios (BI Projections)**





No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Section 9. Regulatory Landscape

## EU Farther Ahead Than US; Big Tech Exposed

Trust and content safety for generative AI needs to be strengthened, and large participants including Snap, Meta, TikTok and Alphabet are well positioned to detect and prevent deep fakes created by its misuse. Doing so should improve brand safety for advertisers and increase conversion rates for ad spending on these platforms. Any potential increase in AI-related regulation also could contribute to greater outlays on securing and encrypting data. Europe is farther ahead than the US on regulating AI, and given the tool's rapid development, we believe that creating a dedicated agency is one of the few workable regulatory models.

## 9.1 EU AI Rules Add to Regulatory Thicket

Apple's delayed rollout of Apple Intelligence in Europe isn't due just to the EU's AI rules, but also a thicket of regulations the EU has imposed on tech in recent years that can make clearance for new product launches onerous and, if done incorrectly, costly. The Digital Markets Act is the primary cause of Apple's delay, while the General Data Protection Regulation (GDPR) has halted Meta, X and Google from training their large-language models on EU user data

The AI Act began subjecting general-purpose AI models like ChatGPT to heightened transparency rules from early 2025 (Fig. 62), posing some obstacles to the rollout of AI in Europe, and likely putting near-term pressure on AI investment and innovation until there's more clarity on the interplay between the AI Act and other regulations.

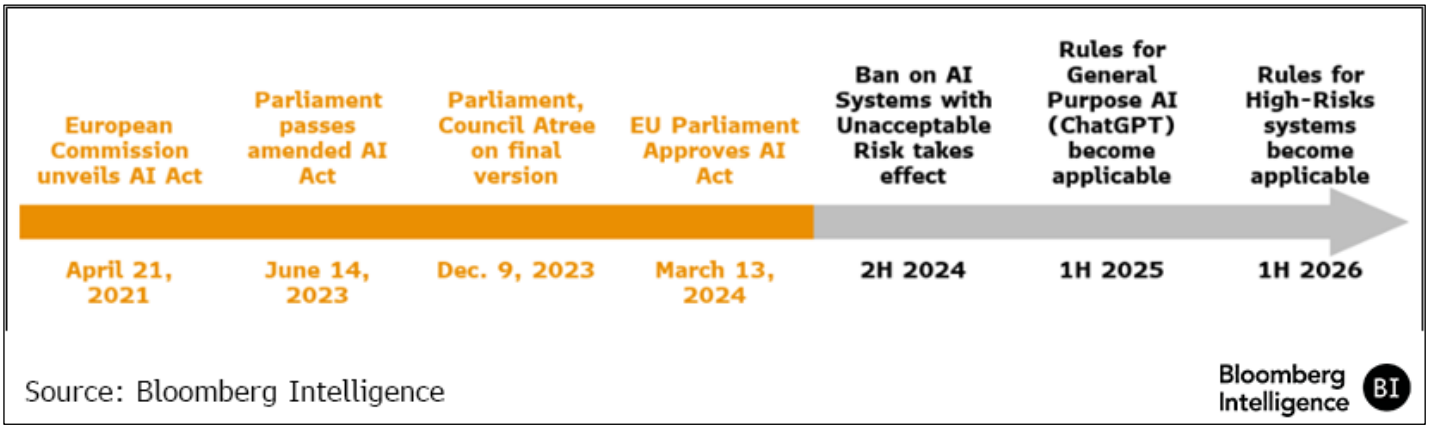
BI **It's doubtful that the US will take aggressive regulatory actions**

The transparency obligations will only apply to general-purpose AI systems that are deemed to pose a "systemic risk." The EU will use training computing power as a basis for bucketing general-purpose AI models into the systemic-risk category, with an initial cutoff of 10 to the 25th power of floating-point operations per second (FLOPs). OpenAI's GPT-4 and Google DeepMind's Gemini are likely to be the first to fall into that category, according to the European Commission. The FLOPs threshold could be adjusted upward or downward over time. Meta has said it won't release its most-powerful AI system in the EU due to regulatory uncertainty. Unless it reverses course, it could avoid systemic-risk designation.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 62: Timeline of EU Act Legislative Process**




The EU's AI Act buckets systems into categories based on the type of risks their applications pose. The rules will outlaw all "unacceptable risk" applications by the end of the year. That includes the use of AI for things such as behavioral manipulation and biometric identification. For "high-risk" applications, which comprise large platforms' recommendation systems, a multistep approach to get approval will take effect in 2026. The European Parliament sought to include general AI systems as high risk, but, in a boon to the industry, that was rejected. Those systems will instead be subject to transparency requirements from next year. Limited-risk applications (such as chatbots) would simply require disclosure, while minimal-risk applications (like the use of a spam filter) wouldn't have any restrictions.

The AI Act could impose fines of up to 7% of annual turnover, exceeding the 4% maximum under the current General Data Protection Regulation. In its first five years of enforcement, the GDPR resulted in cumulative fines of almost €4 billion, with penalties on Meta making up 64% of the total. Amazon accounted for 19% and Google, 5%. We don't believe the threat of penalties would deter near-term investment in generative AI, given the potential to shape the market.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 63: Rules Slowing Rollout of AI Products**

| Firm | Action | Regulation |
|---|---|---|
| Apple | Will not roll out Apple Intelligence in Europe at same time features are offered in other regions | Digital Markets Act |
| Meta | Will not train large-language-model using user-generated content from EU users | General Data Protection Regulation |
| Google | Being probed by data privacy regulators | General Data Protection Regulation |
| Microsoft | Will not train LinkedIn's large-language-model using user-generated content from EU users | General Data Protection Regulation |
| X | Will not train large-language-model using user-generated content from EU users | General Data Protection Regulation |

Source: GDPR Enforcement Tracker, Bloomberg Intelligence

Bloomberg Intelligence BI

Transparency obligations for general-purpose AI models under the EU's AI regulatory package are significantly less onerous than the multistep approval process contemplated for advanced AI systems under an earlier EU proposal. Models that pose "systemic risk" will have additional rules, the most burdensome of which are to assess and mitigate risks. Even that heightened level of scrutiny looks more lax than the pre-release conformity assessment contemplated before.

Obligations to develop and adhere to a copyright policy, and to make available a detailed summary about the content used to train the models will likely fuel more licensing agreements between publishers and developers of large language models.

The AI Act would add to the EU's increasingly complex regulatory framework on tech. In recent years, the bloc has sought to rein in the wanton collection and use of personal data (the GDPR) and to impose obligations on platforms relating to content moderation (the Digital Services Act) and abuse of market power (the Digital Markets Act). The rules can impose substantial financial penalties and operational remedies that can materially alter a business. Enforcement of the GDPR has been fractured, with Ireland taking the mantle for regulating most platforms. The European Commission monitors DSA and DMA compliance. Individual national regulatory bodies likely would lead AI Act enforcement.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

BI **A look at the European Union's recent tech regulations**

**Figure 64: EU's Recent Tech-Focused Regulations**

| Regulation | Compliance Date | Purpose | Mostly Likely Targets | Max Fine |
|---|---|---|---|---|
| General Data Protection Regulation (GDPR) | May 25, 2018 | Provide strict rules on the collection and use of personal data | Meta, Amazon and Google have been the most targeted by GDPR regulators | 4% of annual revenue |
| Digital Services Act (DSA) | August 25, 2023 | Impose new obligations on platforms to enhance moderation of illegal content | EC has identified 17 "Very Large Online Platforms" and two "Very Large Online Search Engines" that were required to comply with the DSA from August 25, 2023 | 6% of annual revenue |
| Digital Markets Act (DMA) | May 2, 2023; March 6, 2024 | Provide new competition rules on large, "gatekeeper" platforms | Gatekeepers include large tech platforms, which will then have to comply by March 2024 | 10% of annual revenue |
| AI Act | 2026 to 2027 | Provide broad safeguards on the development and deployment of AI systems | Potentially any developer or implementer of an AI system, but large tech firms once again likely to be focus given early investments into generative AI models | 7% of annual revenue |

Source: Bloomberg Intelligence

Bloomberg Intelligence BI

## 9.2 Aggressive Regulatory Tack Unlikely in US

If the US adopted an aggressive regulatory approach – which we doubt will happen – it could dent growth of AI products from a range of companies: chipmakers like Micron and Nvidia; cloud infrastructure providers like Amazon and Oracle; software- and development-tool makers like Adobe, IBM and Microsoft; and platforms that use AI for data, search and ad capabilities, like Alphabet and Meta.

The first federal bipartisan bill – the No Section 230 Immunity for AI Act, which would do little – shows how much work remains before serious AI limits land in the US. The measure would clarify that a federal liability shield, Section 230 of the Communications Decency Act, wouldn't apply to AI, but we think courts would only rarely apply the provision to the technology anyway. More notably, the legislation wouldn't create a federal right to sue over AI harms. Nor would it attempt to say what AI is, defining "general artificial intelligence" so broadly that the law could remove Section 230's shield from many existing online platforms. The provision most likely won't become law without being narrowed significantly.

A second bipartisan bill, the National AI Commission Act, strikes us as a sensible first step for Congress, with a solid chance of becoming law. It would create an independent, bipartisan commission with 20 members to study AI's risks and suggest guardrails. The panel would release three reports over two years, shaping regulations. The commission would include industry representatives, which should ease opposition.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

A Senate committee hearing in May generated surprising support for what we see as the most logical – though probably the most disruptive – approach: creating a dedicated federal agency. Congress is incapable of keeping up with AI's rapid development and though a new agency would struggle as well, it at least would have a shot at monitoring the technology and fashioning legal limits. Such a body could also be more focused than an entity like the Federal Trade Commission, which has oversight of all industries. A new agency would upset the apple cart, however, so it likely would face vigorous industry opposition, and companies that might support it probably would push to narrow the body's powers.

Congress also might consider a licensing model, in which AI applications with a broad reach or the potential to cause severe harm can't be put into operation until they've obtained permission from a new regulator. The approach would face fierce industry pushback, including claims that the move would stifle AI innovation and drive it abroad.

BI **Watch for AI to be regulated by existing tech rules**

A less intrusive form of regulation would focus on transparency, requiring that AI products be disclosed and labeled, and might require that regulators or researchers be able to monitor data that's collected or used.

Since Congress is unlikely to reach consensus on a disruptive approach – like creating a new agency or developing a licensing model – we expect AI to be governed by existing laws, even if they're not always a good fit. The FTC could monitor for unfair and deceptive practices, for example, and existing sector-specific laws may limit AI application by industry. AI should thrive in the US under such a light touch. Yet unlike social-media companies, AI users probably won't receive broad protection from lawsuits from Section 230, which raises the specter of legal liability.

## 9.3 Copyright Headwinds Cloud Near-Term Outlook

Licensing deals could be the most straightforward path to resolving copyright disputes between the developers of large language models and the owners of the content upon which those models were trained. Dozens of lawsuits across the US will test whether it was allowable "fair use" to train LLMs on copyrighted content, and, though that's a plausible defense, we believe it's unlikely to protect LLM developers in every case. If the defense fails, damages could reach the tens of billions of dollars. We thus see broad content-licensing arrangements as the most likely path forward. Such proposals may further entrench deep-pocketed early movers in generative AI, like OpenAI, Anthropic, Meta, Amazon and Alphabet.

Dozens of lawsuits brought by newspaper publishers like the New York Times, authors, music publishers and other rights holders, have accused OpenAI and other developers of large-language models of copying without permission. A copyright holder has exclusive rights to reproduce its work, as well as rights to downstream derivative works. A defendant can claim fair use, though it's a determination based on facts that may not develop until trial nears and it could fail. These early cases could test LLM developers' risk appetite and set a precedent on whether licensing training data becomes the norm, which we think is a likely outcome.

A case filed by the New York Times against OpenAI and Microsoft could theoretically result in a damages award of over $300 billion, given copyright damages in the US can be assessed at as much as $150,000 per infringed work and the New York Times complaint alone alleged high volumes of copyrighted content was used to train OpenAI's large language models. The


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

publishers also want the LLMs destroyed, which poses existential risk to OpenAI and significantly endangers Microsoft's investment.

**Figure 65: NYT v OpenAI; UMG v Anthropic Damages Scenarios**

| Comp | $/Work | Potential OpenAI Damages (10 million works | Potenial Anthropic Damages (500 works) |
|---|---|---|---|
| Google Books (low) | $60 | $600 million | $30,000 |
| Statutory Minimum (adjustments) | $200 | $2 billion *Damages Floor* | $100,000 |
| Google Books (high) | $300 | $3 billion | $150,000 |
| Statutory Minimum (no adjustments) | $750 | $7.5 billion | $375,000 |
| Statutory Maximum (non willful) | $30,000 | $300 billion | $15 million |
| Statutory Maximum (willful infringement) | $150,000 | $1.5 trillion *Damages Ceiling* | $75 million |

Note: Google Books comps come from a 2008 proposal by Google to settle copyright infringement claims. The settlement was rejected by a court. Other comps from U.S. Copyright Act that allows for statutory damages ranging from $750 to $30,000 per infringed works. If the infringement is willful, the amount can be increased to $150,000 per work. The court has discretion to reduce the amount to not less than $200 per work, making this the floor for any potential damages award.

Source: Bloomberg Intelligence

Bloomberg Intelligence BI

A defendant sued for copyright infringement in the US can claim "fair use" -- a multifactor balancing test that can negate an infringement claim. It's a powerful defense -- most relevantly used by Google when its Books project was accused of infringing millions of works -- yet Google prevailed on fair use only after a decade of litigation. The overhang was a mild annoyance, but for OpenAI and Anthropic, the litigation may threaten their existence.

Quick resolution is improbable, prolonging a cloud of uncertainty for firms creating and investing in Gen AI systems. Fair use is a highly fact-sensitive determination, making it unlikely to be decided on a dismissal motion. Success on summary judgment is possible, but a recent summary judgment decision on fair use went against the LLM developer --Ross Intelligence--, further underscoring how risky a fair use bid may be.

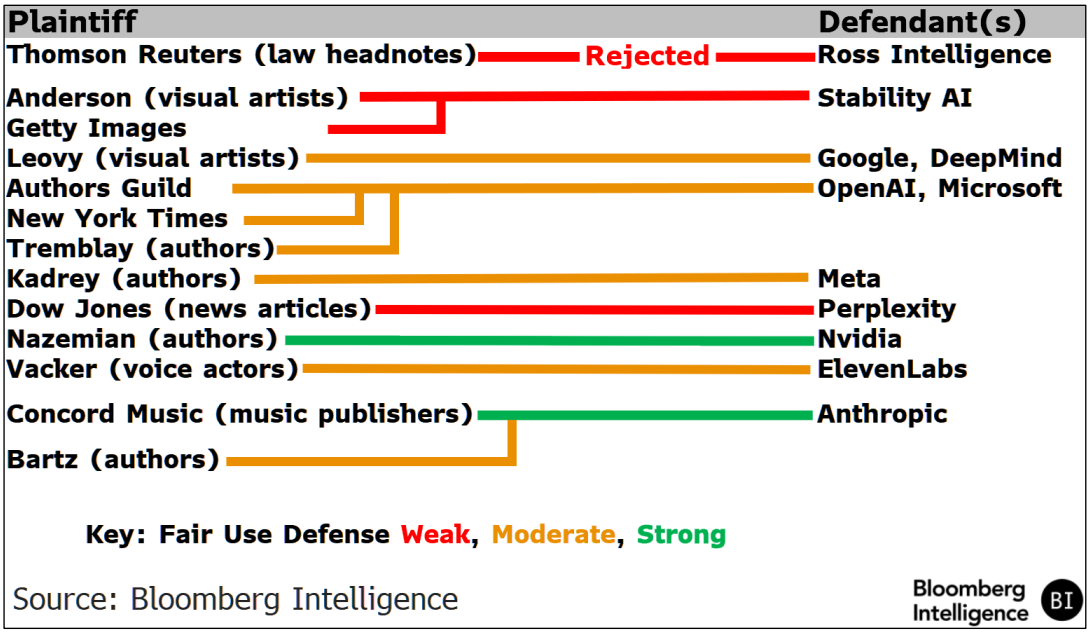
Ross Intelligence's loss on fair use, at the summary judgment stage, could be most damaging to other AI companies that sought to directly compete in the same market as the rights holders whose data were used to train the LLM. Stability AI and Perplexity look most vulnerable to us, as both face lawsuits with relatively strong direct market competition claims.

Fair use will likely be a defense in any case that proceeds to the summary judgment stage. Actions against Meta and Anthropic could reach that point in 2025.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 66: Fair Use in AI Copyright Cases**




Annual royalty payments to owners of content on which large language models are trained could be hundreds of millions of dollars, considered as a whole, and these fees would likely be passed through to end users. LLMs and content owners have struck scores of deals to date, the largest known transaction being a $60 million annual payment to Reddit from Google. Others include a reported $25 million-a-year accord between OpenAI and News Corp., as well as an agreement involving HarperCollins and Microsoft that offers $5,000 per book.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 67: Known Licensing Deals Between LLMs, Content Owners**

| Content Owner | Type | Licensee | Amount |
|---|---|---|---|
| John Wiley & Sons | Academic | Multiple LLM developers | $44 million |
| Informa | Academic | Microsoft, unknown second LLM developer | $75 million |
| HarperCollins | Non-fiction Books | Microsoft | $5,000 per book |
| Associated Press, Axel Springer, Conde Nast, Dotdash Meredith, Financial Times, Future, Guardian, Hearst, News Corp, Schibsted, The Atlantic, Time, Vox Media | Newspapers | OpenAI | News Corp deal alone reportedly cost $25 million/year. Likely aggregate of over $100 million annually. |
| Axel Springer, Financial Times, Hearst, USA Today | Newspapers | Microsoft (Copilot partners) | Unknown |
| Agence France-Press | Newspapers | Mistral | Unknown |
| Time, Der Spiegel, Fortune, Entrepreneur, Texas Tribune, Wordpress | Newspapers /Magazines/ Blogs | Perplexity | Revenue sharing model |
| Reddit | Social Media Content | Google | $60 million/year |
| Shutterstock | Images | OpenAI | Unknown |

Source: Bloomberg Intelligence

Bloomberg Intelligence BI

Mature copyright licensing arrangements between publishers and first-mover generative AI developers could significantly increase barriers to entry for other LLM developers. The models are trained on massive troves of data, and the inclusion of fact-checked news content is helpful in the downstream accuracy of Gen AI output. OpenAI, Anthropic and other first movers allegedly scraped content first, seeking permission only afterward, prompted in part by a slew of lawsuits. OpenAI and Anthropic were on a solid financial footing by the time scrutiny arose over training data, while Alphabet, Microsoft, Meta and others developing LLMs have ample cash flows to absorb licensing costs. New entrants will likely face that scrutiny earlier.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Section 10. ESG Outlook

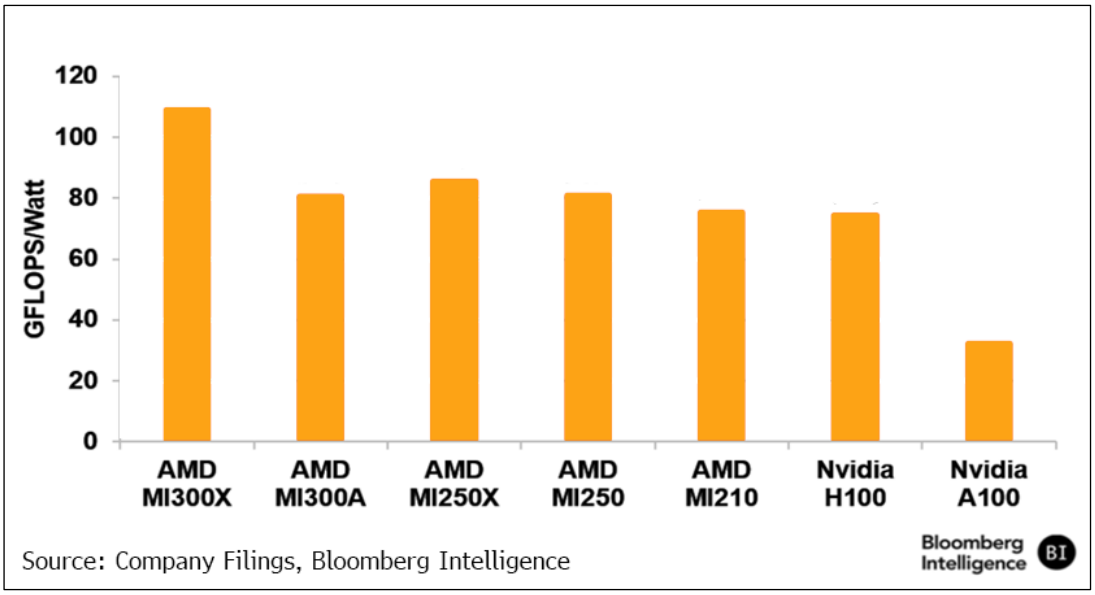
## Reducing Power Consumption; Protecting IP and Privacy

Increased use of graphics processing units for AI inference will significantly boost energy consumption in data-center servers, putting a priority on energy savings to maximize operating efficiency while minimizing power and cooling costs. Meanwhile, about 40% of respondents in a BI survey expressed concern over how generative AI's use of information could breach intellectual-property rights, with privacy another concern.

## 10.1 AMD's Latest Can Challenge Nvidia

That could favor AMD over Nvidia (Fig. 68) as the former's latest MI250X accelerator outperforms Nvidia's H100 by performing a higher peak number of floating-point operations a second per watt. We believe ARMs can continue to gain market share from x86 processors in data centers. Existing internet workloads largely run on x86 architecture, but most generative AI applications will run on chips like GPUs that can conduct massive parallel processing with low power consumption.

**Figure 68: Energy Efficiency of Advanced AI Chips**



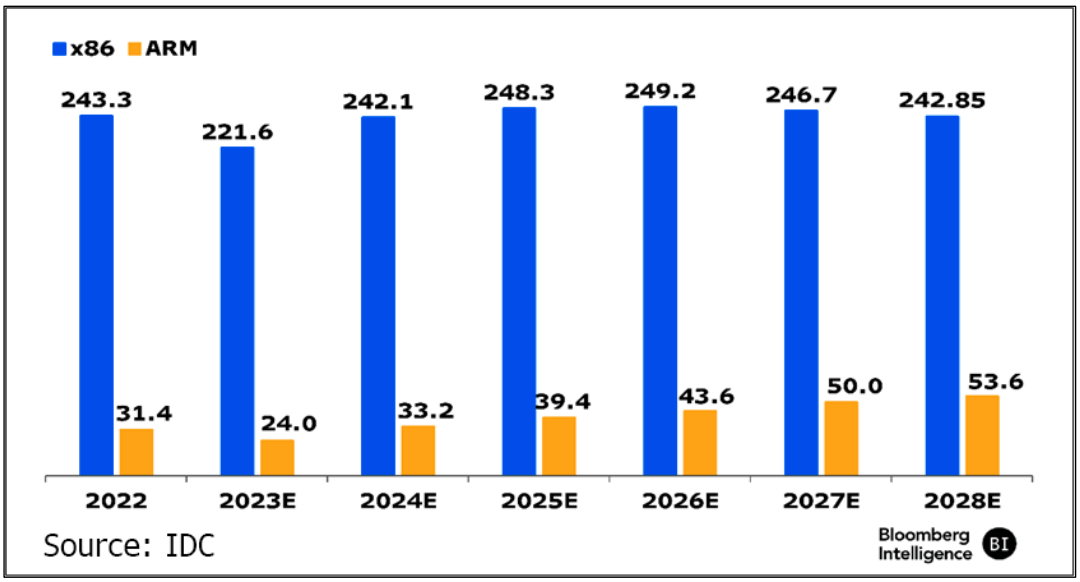

Nvidia's advantage with its CUDA interface is likely to remain for ARM-based processors and may help the company gain market share over x86 processors across data centers, networking and edge devices thanks to the ARM design's power efficiency.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**x86 shipments seen steady, while ARM climbs through 2028**

**Figure 69: x86 vs. ARM Shipment Forecast**




A majority of survey respondents said they still would use gen AI tools despite concerns about intellectual property rights, as long as they provided better results than traditional search functions from Google and other websites. IP concerns will likely decrease over time as companies share more information about how their algorithms are trained using proprietary rather than open web data.

Developers of high-level generative AI systems could retreat from the EU, absent further revisions to the bloc's proposed regulations. Mentions of generative AI and ChatGPT by European companies exploded this year, suggesting a strong desire to employ the new technology. Yet the European Parliament on June 14 adopted rules that would subject developers of generative AI models to additional constraints, such as transparency rules related to the data sets used for training.

Since all AI essentially is driven by data – collecting lots and using machine learning to produce outputs based on it – US regulators could address related harms by limiting what can be collected and used. That would be similar to data-privacy regulations that have been proposed for social-media companies. Given the parallels, expect internet platforms to vigorously lobby against such limits.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Section 11. Performance and Valuation

## AI an Accelerating Tech Theme Facing Pressure

Within the Bloomberg AI Aggregate EW Index (BAIAT Index), based on the BI AI theme universe, pick-and-shovel exposures (hyperscalers, semiconductors and hardware) were the earliest tailwind beneficiaries, though as AI is increasingly integrated, there has been a broadening of the theme's impact across exposures and BAIAET reached its peak EV-to-sales in December 2024.

## 11.1 Performance: Extreme Premiums to Open 2025

A fall from grace thus far in 2025 for our BAIAET Index, down 6%, after what was its top ranking 29% return for full year 2024. The theme remains vulnerable to further macro pressure, with its 89% technology and 62% US weightings leaving it exposed to risk of a deepening correction, as some AI-inflected stocks are still trading at extreme premiums and it remains to be seen if the resiliency of the thematic premium that was established in 2024 can hold amidst broad market volitility. Bright spots are limited within BAIAET though AI stocks in China, up 21%, and Korea (13%), are some.

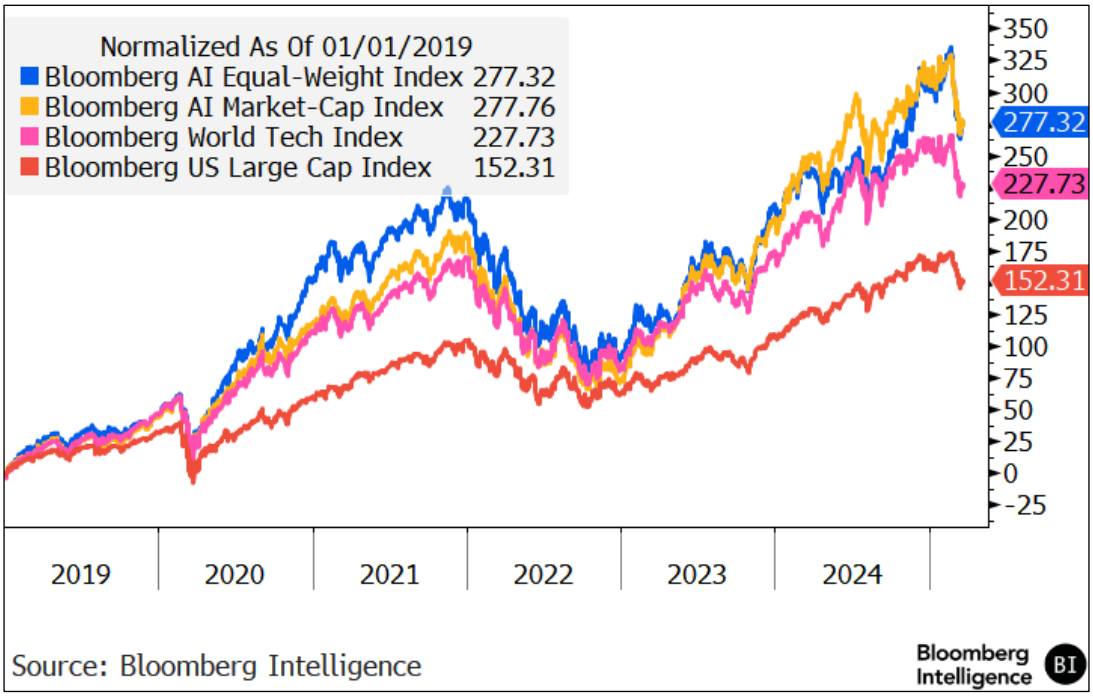
Extended periods of volatility can prove detrimental to investor risk appetite, generally a headwind for high innovation exposed themes, like AI. Though we continue see pockets of resilience in that investor thirst for innovation which suggests short-lived retreats need not translate into longer-term derailment. thematic private-company deal activity is robust, with $26 billion more raised in the first 70 days of 2025 compared with the same period last year. Accelerating tech themes continue to be a draw -- raising $28 billion -- with companies exposed to frontier tech attracting $11 billion vs. just $2.3 billion in the same period last year. AI deal values are up $3.4 billion. Earlier this month, Anthropic's Series E round yielded $3.5 billion, with public companies including Cisco and Fidelity and VCs including Lightspeed Venture Partners and Deer Management all in the mix, pushing its valuation to $61.5 billion. Other frontier tech/AI privates Celestial AI, Reflection AI, Cartesia, Norm Ai and Bria also recently drew capitalOn our 2025 radar: 1.) Investors are shook and volatility reigns as sector and regional rotations abound, but thirst for innovation holds, look for venture capitalists to show affection for the intersection of AI and quantum computing and potentially disruptive energy tech like nuclear fusion. 2) Thematic equities' expansion of premia over non-thematic stocks and the power of equal-weight indexes.

Of BI's 13 accelerating tech themes, AI ranked second to only modern global defense on three-year annualized returns near the end of 2024 and among the top half of those themes on one-year returns. As proliferation of AI solutions persists, business models across industries stand to see disruption making AI arguably the most disruptive theme in our lineup. One of the most challenging differentiations to make when an investment theme gains traction though is point-in-time enthusiasm versus emerging secular growth opportunity. While stock price performance is one of the most turned-to metrics, and AI holds up well broadly on those metrics, for AI, on our radar are fundamental metrics that we see pointing to staying power


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

One of AI's most powerful implications is that its proliferation stands to extend its impact to adjacent themes such as 5G, cloud, cybersecurity and semiconductors. AI and AI-adjacent themes, thus, rank at the top of BI's near-term theme radar.AI and its adjacent themes have robust overlap, with 19 names in three or more of those themes (strong AI-adjacent theme breadth stocks). Using AI as the epicenter, most strong theme breadth stocks are seen in cloud and cybersecurity. Broadcom and Cisco appear in four or more themes, making them unique. Broadcom is in AI, 5G, cloud, cyber and semis; Cisco is in AI, 5G, cloud and cyber.

**Figure 70: Bloomberg AI Index vs. Benchmark Since 2019**




## 11.2 Valuation: AI-Exposed Sectors Look Pricey

Sectors influenced by artificial intelligence are still screening as inordinately expensive -- like US and emerging market tech and US utilities -- even after the market's swoon. The median US tech stock trades at a forward price-to-book ratio around 6.1x, but the relationship between global sector valuations and return on equity estimates imply a 4.8x multiple (20% lower) would be more appropriate. EM tech screens as even pricier, with a median forward price-to-book of 3.2x, while ROE estimates imply a 2.1x multiple (35% lower). Developed market ex-US tech is more reasonably valued, likely given its minimal exposure to the AI supply chain. ROE estimates still imply a median price to book 12% lower than current levels. Developed market ex-US health care and utilities, as well as US utilities, screen as pricey as well.

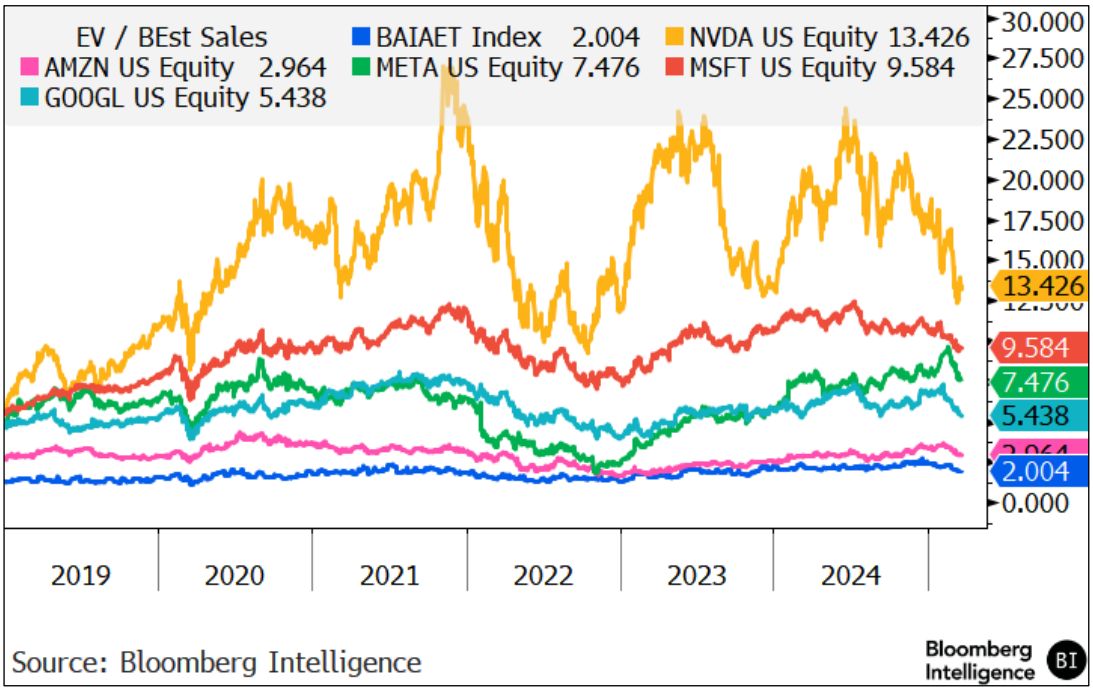
The AI theme, 89% weighted to tech, reached peak EV-to-sales in mid-December, compared with the Bloomberg World Tech Index's (WLSTT) July apex. Today, EV-to-sales are down 26% for BAIAET and 18% for WLSTT from those 2024 peaks, and the near-term holds a potential mix of opportunities and challenges.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

There was significant multiple expansion for specific semiconductor companies in 2024, with Nvidia again leading the way, with the most visible impact on sales growth expectations from AI.

Momentum in top-line estimate revisions could hinge on the pace of product releases such as Microsoft's release of its GitHub and Office copilots. Alphabet has recently released its Gemini LLM and Duet AI copilot while Meta has open-sourced its Llama foundational models to spur adoption. Database and infrastructure software companies such as Oracle, Snowflake, MongoDB and Databricks have continued to ramp up their offerings with vector search capabilities that could stand to benefit from large amounts of data used for training LLMs, which could help drive positive revisions to consensus. However, the benefits that accrue may not be balanced across the players.

**Figure 71: EV-to-Sales: BAIAET, WLSTT Index, Hyperscalers**





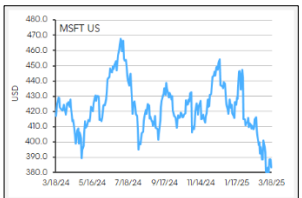

No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Section 12. Company Impacts

## Seeking a Slice of Generative AI's $1.8 Trillion Pie

With a projected $1.8 trillion in spending by 2032, generative AI's effects will ripple through every industry in the technology sector. Here's a look at how some companies are poised to gain over the coming decade.

## 12.1 Microsoft Among Best Positioned Software Makers




| | | |
|---|---|---|
| Market Cap (USD Bn) | | 2875.76 |
| Share Price (USD) | | 383.52 |
| Earnings Per Share* (USD) | | P/E |
| 2024A | 11.80 | 32.50 |
| 2025E | 13.14 | 29.19 |
| 2026E | 15.04 | 25.51 |

* Bloomberg Consensus

**$13 Billion**
Gen AI annual recurring revenue in 2024

**75%**
GitHub Copilot adoption

**Company Outlook:** Microsoft's vast array of software applications makes it a key beneficiary of the growing digital transition as companies upgrade legacy IT systems. A foothold in cloud infrastructure, coupled with its close relationship with OpenAI, puts Microsoft in a strong position to capitalize on rising demand for generative AI.

**AI Impact:** Microsoft is better positioned than most software companies to capitalize on the increased adoption of generative AI, given a first-mover advantage from its tie-up with OpenAI, resulting in $13 billion in generative AI annual recurring revenue in calendar year 2024. It's the first large company to launch AI copilots across its product portfolio, from Office to GitHub. Azure, a cloud-infrastructure product, most likely will be Microsoft's main beneficiary of increased AI demand in the long term. Not only does ChatGPT run on Azure, but Microsoft is also making OpenAI's LLMs available on the platform. Search is another area where we believe Microsoft could gain market share steadily over time.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

## 12.2 Amazon Can Gain From Training, Inference, Creative




| | | |
|---|---|---|
| Market Cap (USD Bn) | | 2055.91 |
| Share Price (USD) | | 192.82 |
| Earnings Per Share* (USD) | | P/E |
| 2024A | 5.53 | 34.87 |
| 2025E | 6.34 | 30.40 |
| 2026E | 7.63 | 25.26 |

* Bloomberg Consensus

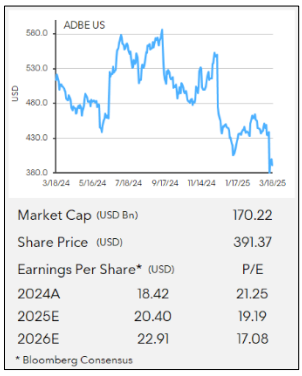
**100-200 Bps**
AWS sales growth lift from Gen AI

**$75 Billion**
Potential 2025 spending to ramp up AI presence

**Company Outlook:** Amazon.com's push for speed, convenience and value, coupled with building momentum across retail, cloud and ads, position it well to deliver on its 1Q plan. Cloud-services growth could accelerate in 2024 with margins also expanding. Improving IT budgets and companies' greater willingness to shift infrastructure to the public cloud remain catalysts for AWS in the longer run. Operating margin may continue to expand on cost cuts, efficiencies and rising contribution from cloud and ads. Amazon's push for pharmacy and grocery are large undertakings that we'll monitor closely

**AI Impact:** AWS should see its fair share of new generative AI workloads from training and inference to creating new applications through its Bedrock offering. Like Microsoft's GitHub, AWS also is offering a generative AI-embedded software development product called CodeWhisperer that significantly shortens the time it takes a developer to code. AWS has the largest market share in cloud infrastructure as a service and more than 100,000 customers using its other AI and machine learning services. Though AWS hasn't combined with OpenAI LLMs yet, it does provide its own foundational models in addition to working closely with other providers, such as Anthropic, Stability AI and AI21.

## 12.3 Adobe to Parlay Base of 70 Million Creative Professionals




| | | |
|---|---|---|
| Market Cap (USD Bn) | | 170.22 |
| Share Price (USD) | | 391.37 |
| Earnings Per Share* (USD) | | P/E |
| 2024A | 18.42 | 21.25 |
| 2025E | 20.40 | 19.19 |
| 2026E | 22.91 | 17.08 |

* Bloomberg Consensus

**$1.4 Billion**
High end of Gen AI sales gain over 2-3 years

**25%-Plus**
DF copilot adoption Year 1

**100 Bps**
Operating margin expansion in 2024

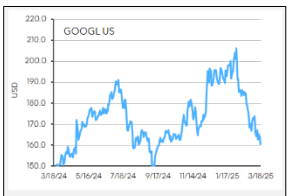
**Company Outlook:** Adobe's solid portfolio of digital products could lead to organic sales growth of 12-15% in constant currency over the next three years, along with non-GAAP operating margin of about 45%. We see the company as among the key beneficiaries of increased spending on digital transformation, given its focus on data insights, digital commerce, marketing and content creation. Generative AI product Firefly could help drive up average revenue per user, as Adobe differentiates itself with built-in copyright safeguards for immediate commercial use, amid increasing competition from providers like OpenAI.

**AI Impact:** Adobe's installed base of around 70 million creative professionals, the highest market share in this category, positions it well to reap benefits from generative AI. The recently launched Firefly creative copilot can substantially reduce the time it takes to create images via text and has already helped create more than a billion visuals through generative fill. The troves of data housed in Adobe's creative cloud suite, which includes Photoshop and Illustrator, make it better positioned than rivals to train its LLMs, and rights to the underlying training content in Adobe Stock can also provide creators with legal peace of mind. Gen AI advancements in digital documents as well as its enterprise front-office software can also aid productivity and could lead to pricing improvements.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

## 12.4 Alphabet Leveraging AI Across Product Suite




| | | |
|---|---|---|
| Market Cap (USD Bn) | | 1993.37 |
| Share Price (USD) | | 160.67 |
| Earnings Per Share* (USD) | | P/E |
| 2024A | 8.04 | 19.98 |
| 2025E | 8.99 | 17.86 |
| 2026E | 10.27 | 15.64 |

* Bloomberg Consensus

**$3-$4 Billion**
Google cloud sales

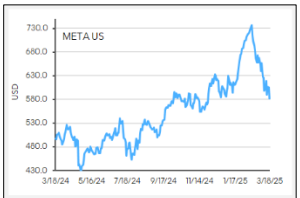
**10-15%**
Search queries to leverage AI

**5-10%**
YouTube ad engagement

**Company Outlook:** Alphabet's improved top-line growth for core search and YouTube segments appear sustainable for the rest of 2024, while demand for generative AI computing bodes well for its Cloud segment profitability. Though Bing-ChatGPT is a risk to the ad business, we believe the company's coming launch of a multimodal large language model and integration of generative AI features into its core search and YouTube products have eased near-term competitive pressure. Network ad sales face pressure amid the removal of cookies, though YouTube ads and subscriptions could see double-digit growth in 2H.

**AI Impact:** Alphabet is exposed to most segments of the generative AI market, including training, inference and digital ads. Gemini is already being used across the company's offerings, including ad targeting algorithms, Google Cloud Vertex AI, and the Google Pixel, the first smartphone with a native generative AI assistant. Cloud AI tools may offer monetization opportunities, while improvements to targeting algorithms can provide a lift to ad pricing. A potential partnership with Apple will further aid Gemini's positioning in the nascent gen AI inferencing market. Though many retailers are rolling out their own enhanced search offerings with longer prompts, the probability of Alphabet losing consumer traffic is relatively low given users' preference for Google's search

## 12.5 Meta's Llama for Recommendations Draw Enterprises




| | | |
|---|---|---|
| Market Cap (USD Bn) | | 1474.97 |
| Share Price (USD) | | 582.36 |
| Earnings Per Share* (USD) | | P/E |
| 2024A | 23.86 | 24.41 |
| 2025E | 25.29 | 23.03 |
| 2026E | 29.01 | 20.07 |

* Bloomberg Consensus

**5-10%**
Effect on engagement, impressions growth

**$1-$2 Billion**
LLM licensing sales bump

**$10 Billion**
Click-to-message ad run rat

**Company Outlook:** Meta may continue to see benefits to user engagement from pivoting to AI-based recommendations, driving stronger impressions growth across its family of apps. The company has leveraged its Llama model to improve ad targeting algorithms while also launching new products like AI Studio and Meta Assistant which can aid monetization efforts. Reels, a high-single-digit contributor to sales, and click-to-messaging ads are both above $10 billion in revenue run rate. User growth will be driven mostly by Instagram and WhatsApp, while Reels could emerge as a contributor to ad pricing this year. Free cash flow may remain squeezed by a slightly higher capex guide and potentially a $20 billion annual loss from Reality Labs in 2024

**AI Impact:** Meta's scale in running its own data-center infrastructure, coupled with large amounts of training data from its family of apps, has allowed it to build its own foundational LLM, Llama, to compete with offerings from OpenAI, Alphabet, and others. Llama's open-source nature may be more attractive for enterprises looking to build their own functionality on top of the model. The pace at which new content is created for social media and virtual- and augmented-reality applications for the metaverse could speed up with generative AI. Meta may also implement personalized shopping assistants to boost user adoption of social commerce, increasing monetization opportunities. Our analysis suggests that the generative AI market may add almost $207 billion in ad spending through 2032 with time spent on platforms, plus ad targeting and personalization.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

## 12.6 OpenAI Supported by Foundational LLM Adoption

**$260 Billion**
Valuation

**$3.7 Billion**
Annual Sales

**Company Outlook:** OpenAI recently raised $6.6 billion for a $157 billion valuation, which was later raised to around $260 billion in another funding round for the company with SoftBank. OpenAI generates about 70% of its $3.7 billion in annual sales run rate from ChatGPT subscriptions, which include 11 million subscribers on its ChatGPT Plus plan and another 1 million on higher-priced plans. OpenAI generates API revenue from both enterprise and individual application programming interface calls, as well as consumption through Azure AI.

**AI Impact:** OpenAI dominant position in generative AI remains aided by its foundational LLM adoption across both enterprise and consumer applications. The company's GPT-4 model came out as the top choice in the Bloomberg Intelligence enterprise CIO survey. OpenAI has also released its o1 model that can reason at the time of inferencing, which other foundational model companies are likely to follow. ChatGPT subscription uptake far exceeds that of rivals Google Gemini and Anthropic Claude aided by higher engagement of the former's chatbot offering.

## 12.7 Anthropic Leverages Partnerships to Post Gains

**$62 Billion**
Valuation

**$500 Million**
LLM License Sales

**Company Outlook:** Anthropic's potential $62 billion valuation in a new private funding round, based on the Information, suggests the company remains among the leading AI native companies. Hyperscaler distribution has so far aided Anthropic's API sales, at about a $500 million run rate given the company's alliances with Amazon and Google.

**AI Impact:** Anthropic's high performance of its Claude LLM across various benchmarks has aided its positioning vs leading rivals including OpenAI, Google and Meta, which all have their own foundational model. Anthropic has partnered with Amazon and Alphabet to deploy its foundational model on public clouds while also sourcing compute for training the next version of its LLM.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Section 13. Methodology

Bloomberg Intelligence's interactive market-sizing model helps provide growth forecasts for generative AI. This model, which will evolve on a regular basis, is still in its early stages, and we have provided Terminal clients an interactive calculator to test scenarios (available at BI INET <GO>). The methodology is based on a bottom-up approach to forecast revenue for areas within hardware, software, digital ads, gaming, IT and business markets. Our forecasts for new segments are anchored to established end-markets that generative AI is likely to disrupt and create new revenue opportunities. The approximate calculations are driven by our assumptions around how generative AI could disrupt these existing end-markets to varying degrees. Figure 72 illustrates the existing end markets with growth assumptions for 2024-28 and 2028-32. Figure 73 shows BI's base-case penetration rates of new AI segments across these end-markets. Figure 74 depicts BI's base-case generative AI revenue forecasts, driven by AI penetration rates.

**Figure 72: Revenue and Growth Forecasts for Existing Technology End-Markets**

| Worldwide Revenue by Technology Group | | | | | |
|---|---|---|---|---|---|
| **(in billions of $)** | **2024** | **'24-'28 CAGR(%)** | **2028E** | **'28-'32 CAGR(%)** | **2032E** |
| **Hardware** | **$1,681** | **6%** | **$2,134** | **10%** | **$3,178** |
| Devices | $1,131 | 3% | $1,248 | 4% | $1,461 |
| Infrastructure | $550 | 13% | $886 | 18% | $1,718 |
| **Software** | **$1,145** | **13%** | **$1,893** | **13%** | **$3,069** |
| Application Development & Deployment | $282 | 20% | $581 | 16% | $1,069 |
| PaaS | $159 | 28% | $426 | 22% | $943 |
| On-premise | $123 | 6% | $155 | -5% | $126 |
| Applications | $591 | 11% | $903 | 10% | $1,320 |
| SaaS | $350 | 16% | $636 | 14% | $1,074 |
| On-premise | $241 | 3% | $267 | -2% | $246 |
| System Infrastructure Software | $272 | 11% | $409 | 14% | $680 |
| IaaS | $135 | 16% | $248 | 22% | $548 |
| On-premise | $137 | 4% | $162 | -5% | $132 |
| **IT Services** | **$884** | **5%** | **$1,085** | **5%** | **$1,341** |
| **Business Services** | **$414** | **5%** | **$502** | **5%** | **$622** |
| **Digital Ad Spending** | **$677** | **10%** | **$993** | **12%** | **$1,571** |
| Search | $276 | 10% | $399 | 10% | $585 |
| Display | $377 | 11% | $563 | 14% | $951 |
| Classified and Other | $24 | 6% | $30 | 4% | $36 |
| **Cybersecurity Spending** | **$117** | **13%** | **$191** | **12%** | **$301** |
| **Gaming Spending** | **$264** | **5%** | **$326** | **8%** | **$444** |
| **Life Sciences Spending** | **$179** | **11%** | **$271** | **11%** | **$412** |
| **Education Technology Spending** | **$140** | **14%** | **$233** | **14%** | **$389** |
| **Total** | **$5,501** | **9%** | **$7,630** | **10%** | **$11,327** |

Source: Bloomberg Intelligence's forecasts based on data from IDC, eMarketer, and Statista

Bloomberg Intelligence BI



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

The hardware market is currently valued at $1.68 trillion, split into devices ($1.131 trillion) and data-center infrastructure ($550 billion), based on IDC data. Bloomberg Intelligence expects new segments for generative AI in this category are AI servers, AI storage, Training Compute, and Networking on the training side, and conversational AI and computer vision products on the inference devices side. The assumptions around the shift in spending to AI servers and storage from traditional servers and storage can be changed in the BI interactive calculator, available on the Terminal. Generative AI infrastructure-as-a-service is how we expect the training compute and storage capacity to be consumed on the cloud. Similarly, for inferencing, conversational AI products and computer vision should emerge as new categories in devices that could be used at home and in cars. For the $1.15 trillion software market, we anticipate new categories to emerge, such as coding copilots, specialized virtual assistants, chatbots and drug discovery software.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Figure 73: Generative AI Penetration Rate (BI Base-Case Assumptions)**

| AI Penetration (%) Base-Case Assumptions | 2024 | 2028E | 2032E |
|---|---|---|---|
| **Hardware** | | | |
| **Devices (Inference)** | | | |
| Computer Vision AI Products | 0.4% | 2.2% | 4.3% |
| Conversational AI Products | 1.1% | 5.7% | 7.8% |
| **Infrastructure (Training)** | | | |
| AI Server | 21.7% | 21.5% | 18.5% |
| AI Storage | 2.7% | 3.9% | 3.4% |
| Training Compute | 0.9% | 2.6% | 2.6% |
| Networking | 1.1% | 2.1% | 2.1% |
| Inference/Fine-Tuning Cloud | 2.7% | 8.6% | 11.2% |
| **Software** | | | |
| **Coding, DevOps and Generative AI Workflows Software** | 0.5% | 4.0% | 7.0% |
| Copilot Spending | 0.4% | 2.4% | 3.5% |
| Customer Service/Contract Review AI Agents | 0.5% | 10.0% | 20.0% |
| **Specialized Generative AI Assistants Software** | | | |
| E-Commerce | 0.1% | 0.4% | 0.9% |
| Social Media | 0.1% | 0.4% | 1.0% |
| **Generative AI Workload Infrastructure Software** | 1.0% | 5.0% | 12.0% |
| **Drug Discovery Software** | 0.1% | 5.0% | 10.0% |
| **Cybersecurity Spending** | 0.5% | 6.0% | 8.0% |
| Copilot Spending | 0.1% | 2.8% | 4.7% |
| Data Protection | 0.5% | 3.2% | 3.3% |
| **Education Spending** | 1.0% | 3.0% | 6.0% |
| Copilot Spending | 0.9% | 2.7% | 5.4% |
| Educational Content Creation | 0.1% | 0.3% | 0.6% |
| **Image/Video Generation Tools** | 0.1% | 0.5% | 1.0% |
| **Film Production/Music Generation Tools** | 0.0% | 0.1% | 0.2% |
| **Gaming Spending** | | | |
| **Virtual Goods** | 0.4% | 3.0% | 6.0% |
| **Game Design Software** | 0.6% | 5.0% | 10.0% |
| **IT Services** | 1.0% | 4.0% | 6.0% |
| **Business Services** | 0.1% | 3.0% | 5.0% |
| **Digital Ad Spending** | | | |
| **Search** | 2.0% | 7.0% | 12.0% |
| **Videos** | 1.0% | 5.5% | 10.6% |
| **Messaging** | 0.4% | 2.0% | 4.0% |

Source: Bloomberg Intelligence

Bloomberg Intelligence BI



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

## 13.1 BI's Market-Sizing Conclusions

The 2024-32 forecast scenario for the generative AI market is based on the size of the end-markets and penetration. The CAGR assumptions for both-end markets and generative AI impact can be modified to come up with your own scenarios. For example, we assume in our base case that the data-center market ("Infrastructure" in Figure 72) is likely to expand at a 12% CAGR from 2024-28 and 18% from 2028-32. The training segment is further comprised of AI Server, AI Storage, Training Compute, Networking, and LLM Licensing. Inferencing includes devices and software applications, such as Chatbots, Copilots, etc.

**Figure 74: Generative AI Revenue Base-Case Forecast by Technology Segment**

| (in milllions of $) | 2024 | 2028E | 2032E | Implied 8 yr. CAGR (%) |
|---|---|---|---|---|
| **Infrastructure (Training)** | **$150,179** | **$325,585** | **$579,526** | **18%** |
| AI Servers | $119,325 | $190,373 | $317,419 | 13% |
| AI Storage | $14,916 | $34,267 | $59,055 | 19% |
| Training Compute (Cloud Workloads) | $4,972 | $22,845 | $44,291 | 31% |
| Networking | $5,966 | $19,037 | $36,909 | 26% |
| LLM Licensing Revenue | $5,000 | $59,063 | $121,852 | 49% |
| **Devices and Applications (Inference)** | **$38,655** | **$285,075** | **$734,732** | **44%** |
| Inference/Fine-Tuning (Cloud Workloads) | $14,916 | $76,149 | $191,928 | 38% |
| Drug Discovery Software | $90 | $13,571 | $41,203 | 115% |
| Computer Vision AI Products | $5,000 | $27,147 | $63,517 | 37% |
| Conversational AI Products | $13,001 | $70,583 | $114,330 | 31% |
| Chatbots/AI Agents | $4,982 | $90,809 | $311,497 | 68% |
| Customer Service / Contract Review AI Ag | $1,411 | $58,084 | $213,862 | 87% |
| Cybersecurity Copilots | $58 | $5,351 | $14,133 | 99% |
| Coding and Devops Copilots | $1,128 | $13,940 | $37,426 | 55% |
| Educational Copilots | $1,262 | $6,303 | $20,995 | 42% |
| E-Commerce Copilots | $532 | $3,340 | $11,880 | 47% |
| Social Media Chatbots | $591 | $3,791 | $13,200 | 47% |
| Customer Service / Contract Review AI Ag | $1,411 | $58,084 | $213,862 | 87% |
| Educational Content Creation | $140 | $700 | $2,333 | 42% |
| Data Protection | $526 | $6,115 | $9,923 | 44% |
| **Generative AI Driven Ad Spending** | **$10,799** | **$70,179** | **$208,987** | **45%** |
| Search | $5,522 | $27,959 | $70,174 | 37% |
| Videos | $3,770 | $30,962 | $100,783 | 51% |
| Messaging | $1,508 | $11,259 | $38,031 | 50% |
| **IT Services** | **$8,838** | **$43,416** | **$80,443** | **32%** |
| **Business Services** | **$207** | **$15,063** | **$31,097** | **87%** |
| **Workload Monitoring Software** | **$2,719** | **$20,459** | **$81,603** | **53%** |
| **Entertainment/Media** | **$3,353** | **$31,344** | **$86,883** | **50%** |
| Gaming | $2,644 | $26,109 | $71,042 | 51% |
| Virtual Goods | $1,058 | $9,791 | $26,641 | 50% |
| Game Design | $1,586 | $16,318 | $44,401 | 52% |
| Image/Video Generation Tools | $591 | $4,513 | $13,200 | 47% |
| Film Production/Music Generation Tools | $118 | $722 | $2,640 | 47% |
| **Total** | **$214,752** | **$791,121** | **$1,803,271** | **30%** |

Source: Bloomberg Intelligence

Bloomberg Intelligence BI



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Section 14. Glossary of Terms

## These can help decipher highly technical elements:

**Advanced RISC Machines (ARM):** A processor architecture based on 32-bit reduced instruction set computer.

**AI Assistants:** A software agent that can perform tasks for a user based on input such as commands or questions. Think Siri or Cortana.

**AI Server:** Computers used for AI inferencing and training.

**AI Storage:** Often a software-as-a-service application that performs analysis in a public cloud.

**ChatGPT:** A free chatbot that can answer just about any question.

**Conversational User Interface:** Allows people to interact with software, apps and bots as they would another human being. Amazon's Alexa is an example.

**Computer Vision:** Enables computers and systems to derive meaningful information from digital information, videos and other visual inputs, then act or make recommendations based on that information.

**Corpus:** Literally "body," this is the collection of billions of data points used to train a large language model.

**CPU:** Central Processing Unit. Basically, the semiconductor chip that is the essential logic circuitry in a hardware system.

**Edge:** Deployment of computing and storage resources closer to where data is produced.

**Ethernet:** Technology to connect devices in a local area network (LAN) or wide area network (WAN). Slower than InfiniBand.

**Generative AI:** Uses algorithms, such as ChatGPT, to create new content including audio, code, images, text and videos.

**GPU:** Graphics Processing Unit. A specialized circuit for image and video display.

**Hallucination:** A response/output from a large language model that is irrelevant or incorrect

**Inference:** The process of reasoning and making decisions based on available information or data. It follows training to derive new knowledge or conclusions from existing data.

**InfiniBand Network:** A high-performance, low-latency way to facilitate high-speed communications. Faster than Ethernet

**IaaS:** Infrastructure as a Service, a business model that offers computing, storage and networking resources on demand.

**Large Language Models:** Deep learning algorithms that can recognize, summarize, translate, predict, and generate content using massive datasets.

**Machine Learning:** The use and development of computer systems that learn and adapt without following explicit instructions, by using algorithms and statistical models to analyze and draw inferences from patterns in data.

**Neural Network:** A method that teaches computers to process data the way the brain does. It's a type of machine-learning that uses interconnected nodes in a layered structure.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**ODM:** Original Design Manufacturer.

**PaaS:** Platform as a Service. A type of cloud computing service model that offers a flexible, scalable platform to develop, deploy, run and manage apps.

**Personalization:** Uses AI and machine learning to analyze a user's data, helping to understand their needs and tailor their experience accordingly.

**Retrieval-Augmented Generation (RAG):** A method by which a company can supply its proprietary data to a model that can append relevant information to a user query prior to feeding it into a LLM

**Training:** The process of teaching an AI system to perceive, interpret and learn from data. That way, the AI will later be capable of inferencing—making decisions based on information it's provided.


No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Bloomberg Intelligence Research Coverage

**Bloomberg Editorial and Research:**
**John Micklethwait**, Editor-in-Chief; **Reto Gregori**, Deputy Editor-in-Chief

## Research Management

**David Dwyer**, Global Director of Research
**Drew Jones**, Deputy Global Director of Research
**Sam Fazeli**, Director of Global Industry Research
**Alison Williams**, Director of Global Strategy Research
**Paul Gulberg**, Director of Americas Industry Research
**Catherine Lim**, Director of APAC Industry Research
**Sue Munden**, Director of EMEA Industry Research
**Frank Jacobs**, Global Chief Operating Officer

## Content Management

**Tim Craighead**, Global Chief Content Officer
**Karima Fenaoui**, Research Content Officer, Communities & EM
**John Lee**, Research Content Officer, APAC
**Renato Prieto**, Research Content Officer, FICC
**Roger Thomson**, Research Content Officer, Americas
**Rod Turnbull**, Research Content Officer, EMEA
**Mariam Traore**, Research Digital Content Specialist
**Matthew Bloxham**, Global Head of Alternative Data & Analytics
**Brian Egger**, Global Head of Financial Modeling
**Donna Weston**, Co-head Global Research Editorial
**Douglas Zehr**, Co-head Global Research Editorial

## Equity Strategy

**Gina Martin Adams, Director of Equity Strategy Research**

### Markets

**Christopher Cain**, Quantitative Analysis, US
**Michael Casper**, Small Caps and Sectors, US
**Marvin Chen**, China and North Asia
**Nitin Chanduka**, Emerging Markets, India
**Laurent Douillet**, Europe
**Anthony Feld**, Technical Analysis, Global
**Kumar Gautam**, Emerging Markets, Global
**Gina Martin Adams**, Global
**Wendy Soong**, Americas
**Sufianti Sufianti**, Emerging Markets, ASEAN
**Gillian Wolff**, Global

### Funds

**Eric Balchunas**, Exchange Traded Funds, Global
**David Cohne**, Mutual Funds, Global
**Henry Jim**, Exchange Traded Funds, Europe
**Athanasios Psarofagis**, Exchange Traded Funds, Americas
**James Seyffart**, Exchange Traded Funds, Americas
**Rebecca Sin**, Exchange Traded Funds, APAC

### Thematic

**Breanne Dougherty,** Global
**Andrew Silverman**, Tax Policy and Corporate Actions, Global
**Shirley Wong**, APAC
**Lu Yeung**, Accounting, Global

## FICC Strategy

**Noel Hebert, Director of FICC Strategy Research**

### Markets

**Erica Adelberg**, MBS, Americas
**Negisa Balluku**, Litigation-Bankruptcy, Americas
**Mahesh Bhimalingam**, Credit Strategy, EMEA
**Philip Brendel**, Distressed Debt, Americas
**Rod Chadehumbe**, ABS, Americas
**Sam Geier**, Credit Strategy, Americas
**Noel Hebert**, Credit Strategy, Americas
**Eric Kazatsky**, Municipals, Americas
**Mike McGlone**, Commodity Strategy, Global
**Tanvir Sandhu**, Derivatives, Global
**Damian Sassower**, Emerging Markets, Global
**Timothy Tan**, Credit Strategy, APAC

### FX and Rates

**Audrey Childe-Freeman**, FX, G-10
**Stephen Chiu**, FX and Rates, APAC
**Ira Jersey**, Rates,US
**Davison Santana,** FX and Rates, LatAm
**Sergei Voloboev**, FX, Emerging Markets, Global
**Huw Worthington**, Rates, EMEA

## ESG Research

**Adeline Diab**, Director of ESG Strategy Research

### Strategy

**Manish Bangard**, Americas
**Adeline Diab**, Global
**Ortis Fan**, APAC
**Yasutake Homma**, Japan
**Rahul Mahtani**, Quantitative Analysis, EMEA
**Christopher Ratti**, Fixed Income, Global

**Eric Kane**, Director of ESG Company Research

### Company and Industry

**Shaheen Contractor**, Industry, Global
**Rob Du Boff**, Governance, Global
**Gail Glazerman**, Integration, Americas
**Eric Kane**, Global
**Grace Osborne**, Integration, EMEA
**Andrew John Stevenson**, Climate, Global
**Conrad Tan**, Integration, APAC

## Market Structure Research

**Larry Tabb, Director of Market Structure Research**

**Jackson Gutenplan**, Equities, Americas
**Brian Meehan**, Fixed Income, Americas
**Nicholas Phillips**, Equities, EMEA
**Larry Tabb**, Equities and Fixed Income, Global



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Credit Research

**Joel Levington, Director of Credit Research**

## Americas

**Himanshu Bakshi**, Consumer Finance, Banking, Global
**Mike Campellone**, Specialty Apparel, Consumer Hardlines, Global
**Cecilia Chan**, Gaming Lodging & Restaurants, Internet Media, China
**Jean-Yves Coupin**, Health Care, Corporate Bonds, Americas
**Spencer Cutter**, Oil & Gas, Global
**Stephen Flynn**, Entertainment, Cable & Satellite, Americas
**Matthew Geudtner**, Aerospace, Global, Machinery, Americas
**David Havens**, Consumer Finance, Investment Mgmt, Americas
**Mike Holland**, Hospitals, Specialty-Generic Pharma, Americas
**Julie Hung**, Packaged Food, Beverages, Americas
**Arnold Kakuda**, Investment Banking, Americas
**Joel Levington**, Automobiles, Global, Industrials, Americas
**Jody Lurie**, Gaming Lodging & Restaurants, Americas
**Robert Schiffman**, Hardware & Storage, Internet Media, Global

## Asia

**Andrew Chan**, Real Estate, Infrastructure, China
**Sharon Chen**, Telecom Carriers, Infrastructure, India
**Pri De Silva**, Banking, Aerospace & Defense, APAC
**Daniel Fan**, Real Estate, China
**Rena Kwok**, Banking, APAC
**Mary Ellen Olson**, Metals & Mining, Oil & Gas, APAC

## EMEA

**Tolu Alamutu**, Real Estate, Banking, EMEA
**Ruben Benavides**, Banking, Europe
**Aidan Cheslin**, Telecom Carriers, EMEA
**Jeroen Julius**, Banking, EMEA
**Stephane Kovatchev**, Industrials, Construction, EMEA, Americas
**Paul Vickars**, Electric Utilities, EMEA, Oil & Gas, Global

# Industry Research

## Consumer

### Consumer Products & Services

**Brian Egger**, Gaming & Lodging, Americas
**Angela HanLee**, Gaming & Hospitality, APAC
**Drew Reading**, Homebuilders, Americas
**Deborah Aitken**, Luxury, Personal Care Products, Global
**Michael Halen**, Restaurants, Americas
**Conroy Gaynor**, Travel & Leisure, EMEA

### Retail & Wholesale

**Lindsay Dutch**, Consumer Hardlines, Retail REIT, Americas
**Tatiana Lisitsina**, Consumer Hardlines, Online Apparel, EMEA
**Poonam Goyal**, E-Commerce, Specialty Apparel Stores, Americas
**Catherine Lim,** Consumer Goods, E-Commerce, APAC
**Charles Allen**, Retail Staples & Wholesale, Specialty Apparel, EMEA
**Mary Ross Gilbert**, Specialty Apparel Stores, Americas
**Abigail Gilmartin**, Athleisure & Footwear, Americas

### Consumer Staples

**Jennifer Bartashus**, Packaged Food and Retail Staples, Americas
**Duncan Fox**, Beverages, Packaged Food, EMEA
**Kenneth Shea**, Beverages, Tobacco & Cannabis, Americas
**Diana Gomes**, Consumer Health, Household Products, Global
**Lisa Lee**, Consumer Goods, Health Care, APAC
**Ada Li,** Consumer Goods, Online Health Care, APAC
**Diana Rosero-Pena**, Packaged Food, Retail Staples, Americas
**Lea El-Hage**, Retail Staples, Australia

## Energy

**Will Hares**, Energy Sector Head
**Brett Gibbs**, Biofuels, EMEA
**Talon Custer**, Oil & Gas, Americas
**Henik Fung**, Oil & Gas, Gas Utilities, APAC
**Will Hares,** Oil & Gas, EMEA
**Scott J. Levine**, Oil & Gas, Industrials, Americas
**Vincent G. Piazza**, Oil & Gas, Americas
**Salih Yilmaz**, Oil & Gas, EMEA
**Rob Barnett**, Solar Energy Equipment, Americas, EMEA

## Utilities

**Patricio Alvarez**, Electric Utilities, Gas Utilities, EMEA
**Nikki Hsu**, Electric Utilities, Americas
**Kelvin Ng**, Utilities, Coal, APAC
**Gabriela Privetera**, Electric Utilities, Americas

## Financials

### Financial Services

**Ben Elliott**, Consumer Finance, Americas
**Edmond Christou**, Financials, Middle East
**Diksha Gera,** Global Payments and Fintech, Americas
**Salome Skhirtladze**, Financials, Middle East

## Industrials

**Steve Man**, Director of Industrial Research

### Automotive

**Joanna Chen**, Automobiles, APAC
**Gillian Davis**, Automobiles, EMEA
**Michael Dean**, Automobiles, EMEA
**Steve Man**, Automobiles, Americas
**Tatsuo Yoshida**, Automobiles, Japan

### Industrial & Industrial Services

**George Ferguson**, Aerospace & Defense, Global
**Will Lee**, Aerospace & Defense, Americas
**Wayne Sanders**, Defense, Americas
**Stuart Gordon**, Business Services, Europe
**Christopher Ciolino**, Machinery, Industrials, Americas
**Christina Feehery**, Industrials, Americas
**Takeshi Kitaura**, Industrials, Japan
**Mustafa Okur**, Electrical Equipment, Industrials, Americas
**Bhawin Thakker**, Industrials, EMEA
**Karen Ubelhart**, Industrials, Machinery, Americas
**Omid Vaziri**, Industrials, EMEA
**Denise Wong**, Infrastructure, APAC

### Transportation

**Tim Bacchus**, Airlines, APAC
**Francois Duflot**, Airlines, Americas
**George Ferguson**, Airlines, Global
**Conroy Gaynor**, Airlines, EMEA
**Lee Klaskow,** Freight Transportation & Logistics, Global
**Kenneth Loh**, Marine Shipping, Logistics Services, APAC

## Materials

**Jason Miner**, Agriculture Sector Head
**Grant Sporre**, Metals and Mining Sector Head

### Chemicals

**Daniel Cole**, Agricultural Chemicals, Americas
**Alexis Maxwell**, Agricultural Chemicals, Canada
**Alvin Tai**, Agriculture, Malaysia, EMEA
**Jason Miner**, Agriculture and Chemicals, Americas
**Sean Gilmartin**, Specialty Chemicals, Americas
**Vivien Zheng**, Speciality Chemicals, APAC

### Metals & Mining

**Richard Bourke**, Basic Materials, Americas
**Michelle Leung**, Metals & Mining, China, Japan
**Emmanuel Munjeri**, Metals & Mining, South Africa
**Alon Olsha**, Metals & Mining, EMEA
**Grant Sporre**, Metals & Mining, EMEA

### Construction Materials

**Sonia Baldeira**, Construction, Building Materials, EMEA



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Financial Services (Cont'd)**

**Sarah Jane Mahmud**, Banking, Market Structure, ASEAN and India
**Alison Williams**, Investment Banking, Global
**Sharnie Wong**, Investment Banking, Exchanges, APAC
**Paul Gulberg**, Investment Management, Exchanges, Banks Americas
**Ethan Kaye**, Investment Management, Americas
**Neil Sipes**, Investment Management, Investment Banking, Americas
**Hideyasu Ban**, Financial Services, Japan
**Matt Ingram**, Financial Services, Australia and Korea

**Banking**

**Francis Chan**, Banking & Fintech, China & Hong Kong
**Herman Chan**, Banking, Americas
**Tomasz Noetzel**, Banking, EMEA
**Philip Richards**, Banking, EMEA
**Lento Tang**, Banking, EMEA
**Maryana Vartsaba**, Banking, EMEA

**Insurance**

**Steven Lam**, Insurance, APAC
**Jeffrey Flynn**, Life Insurance, Americas
**Kevin Ryan**, Life Insurance, P&C Insurance, EMEA
**Charles Graham**, P&C Insurance, Life Insurance, EMEA
**Matthew Palazola**, P&C Insurance, Americas

**Real Estate**

**Jack Baxter**, Real Estate, Australia
**Ken Foong**, Real Estate, Singapore
**Iwona Hovenko**, Real Estate, Business Services, EMEA
**Kristy Hung**, Real Estate, China
**Jeffrey Langbaum**, Residential REIT, Office REIT, Americas
**Sue Munden**, Real Estate, EMEA
**Patrick Wong**, Real Estate, APAC

**Health Care**

**Aude Gerspacher**, Director of Health Care Research

**Biotech & Pharma**

**Sam Fazeli**, Biotech, Global
**Grace Guo**, Biotech, Pharma, Americas
**Jean Rivera Irizarry**, Biotech, Pharma, Americas
**Jamie Maarten**, Biotech, China
**Max Nisen**, Biotech, Americas
**Michael Shah**, Biotech, Specialty-Generic Pharma, EMEA
**Leslie Yang**, Biotech, China
**John Murphy**, Large Pharma, Biotech, Americas, EMEA
**Aude Gerspacher**, Pharma, Biotech, Americas
**Justin Kim,** Biotech, Pharma, Americas
**Ann-Hunter Van Kirk**, Specialty-Generic Pharma, Americas
**Glen Losev**, Hospitals, Managed Care, Americas
**Jonathan Palmer**, Medical Devices, Supply Chain, Americas
**Matt Henriksson**, Medical Equipment & Devices, Americas
**Kevin Kouam**, Building Materials, Global

**Technology**

**Mandeep Singh**, Director of Technology Research

**Hardware**

**Woo Jin Ho**, Hardware & Networking, Americas
**Ken Hui**, Semiconductors, Europe
**Charles Shum**, Semiconductors, APAC
**Jake Silverman**, Logic ICs, Americas
**Kunjan Sobhani**, Logic ICs, Americas
**Steven Tseng**, EMS/ODM, Consumer Electronics, APAC
**Masahiro Wakasugi**, Semiconductors, EMS/ODM, Global

**Software**

**Tamlin Bason**, IT Services, Americas, EMEA
**Robert Lea**, Internet Media, Application Software, China
**Nathan Naidu**, Entertainment Content, Internet Media, Japan
**Niraj Patel**, Application Software, Americas
**Sunil Rajgopal**, Application Infrastructure Software, Americas
**Anurag Rana**, Application Software, IT Services, Americas
**Mandeep Singh**, Software, Internet, Hardware/Semis, Americas

**Communications**

**Media**

**Matthew Bloxham**, Media, Advertising, Telecom, EMEA
**Geetha Ranganathan**, Entertainment, Cable, Advertising, Americas
**Tom Ward**, Media, Advertising, Telecom Carriers, EMEA

**Telecommunications**

**John Butler**, Telecom & Towers, Infrastructure Software, Americas
**John Davies**, Telecom Carriers and Media, EMEA
**Erhan Gurses**, Telecom Carriers, EMEA
**Marvin Lo**, Telecom Carriers, APAC
**Chris Muckensturm**, Telecom Carriers, ASEAN

**Litigation and Policy**

**Nathan Dean**, Financials Policy, Americas
**Holly Froum**, Consumer, Industrials Litigation and Policy, Americas
**Josephine Garban,** Health Care Patent Litigation, Americas
**Jennifer Rie**, Antitrust Litigation and Policy, Americas
**Matthew Schettenhelm**, TMT Litigation and Policy, Americas
**Elliott Stein**, Financials Litigation, Americas
**Justin Teresi**, Antitrust Litigation and Policy, Americas
**Duane Wright**, Health Care Policy, Americas



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Copyright and Disclaimer

## Copyright

© Bloomberg Finance L.P. 2025. This publication is the copyright of Bloomberg Finance L.P. No portion of this document may be photocopied, reproduced, scanned into an electronic system or transmitted, forwarded or distributed in any way without prior consent of Bloomberg Finance L.P.

## Disclaimer

The data included in these materials are for illustrative purposes only. The BLOOMBERG TERMINAL service and Bloomberg data products (the "Services") are owned and distributed by Bloomberg Finance L.P. ("BFLP") except (i) in Argentina, Australia and certain jurisdictions in the Pacific islands, Bermuda, China, India, Japan, Korea and New Zealand, where Bloomberg L.P. and its subsidiaries ("BLP") distribute these products, and (ii) in Singapore and the jurisdictions serviced by Bloomberg's Singapore office, where a subsidiary of BFLP distributes these products. BLP provides BFLP and its subsidiaries with global marketing and operational support and service. Certain features, functions, products and services are available only to sophisticated investors and only where permitted. BFLP, BLP and their affiliates do not guarantee the accuracy of prices or other information in the Services. Nothing in the Services shall constitute or be construed as an offering of financial instruments by BFLP, BLP or their affiliates, or as investment advice or recommendations by BFLP, BLP or their affiliates of an investment strategy or whether or not to "buy," "sell," or "hold" an investment. Information available via the Services should not be considered as information sufficient upon which to base an investment decision. The following are trademarks and service marks of BFLP, a Delaware limited partnership, or its subsidiaries: BLOOMBERG, BLOOMBERG ANYWHERE, BLOOMBERG MARKETS, BLOOMBERG NEWS, BLOOMBERG PROFESSIONAL, BLOOMBERG TERMINAL and BLOOMBERG.COM. Absence of any trademark or service mark from this list does not waive Bloomberg's intellectual property rights in that name, mark or logo. All rights reserved. © 2025 Bloomberg.

Bloomberg Intelligence is a service provided by Bloomberg Finance L.P. and its affiliates. ("Bloomberg"). Bloomberg is not an officially recognized credit rating agency in any jurisdiction, and customers should not use or rely on Bloomberg Intelligence to comply with applicable laws or regulations that prescribe the use of ratings issued by accredited or otherwise recognized credit rating agencies. Bloomberg Intelligence Credit and Company research may not be available in certain jurisdictions.

Bloomberg Intelligence shall not constitute, nor be construed as, investment advice or investment recommendations (i.e., recommendations as to whether or not to "buy," "sell," "hold," or to enter or not to enter into any other transaction involving any specific interest) or a recommendation as to an investment or other strategy. No aspect of the Bloomberg Intelligence function is based on the consideration of a customer's individual circumstances. Bloomberg Intelligence should not be considered as information sufficient upon which to base an investment decision. Customers should determine on their own whether they agree with Bloomberg Intelligence. Bloomberg Intelligence should not be construed as tax or accounting advice or as a service designed to facilitate any Bloomberg Intelligence customer's compliance with its tax, accounting, or other legal obligations. Bloomberg believes that the information it uses in Bloomberg Intelligence comes from reliable sources, but does not guarantee the accuracy of information contained in Bloomberg Intelligence. Employees involved in Bloomberg Intelligence may hold positions in the securities analyzed or discussed on Bloomberg Intelligence.

Bloomberg makes no claims or representations, or provides any assurances, about the sustainability characteristics, profile or data points of any underlying issuers, products or services, and users should make their own determination on such issues.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# About Bloomberg Intelligence

## Your go-to resource for making better investment decisions, faster.

Bloomberg Intelligence (BI) research delivers an independent perspective providing interactive data and research across industries and global markets, plus insights into company fundamentals. The BI, team of 500 research professionals is here to help clients make more informed decisions in the rapidly moving investment landscape.

BI's coverage spans all major global markets, more than 135 industries and 2,000 companies, while considering multiple strategic, equity and credit perspectives. In addition, BI has dedicated teams focused on analyzing the impact of government policy, litigation and ESG.

BI is also a leading Terminal resource for interactive data. Aggregated, from proprietary Bloomberg sources and 600 independent data contributors, the unique combination of data and research is organized to allow clients to more quickly understand trends impacting the markets and the underlying securities.

Bloomberg Intelligence is available exclusively for Bloomberg Terminal® subscribers, available on the Terminal and the Bloomberg Professional App.



No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

**Take the next step.**

For additional information,

press the <HELP> key twice on the Bloomberg Terminal®.

**Beijing**
+86 10 6649 7500

**Hong Kong**
+852 2977 6000

**New York**
+1 212 318 2000

**Singapore**
+65 6212 1000

**Dubai**
+971 4 3641000

**London**
+44 20 7330 7500

**San Francisco**
+1 415 912 2960

**Sydney**
+61 2 9777 86 00

**Frankfurt**
+49 69 92041210

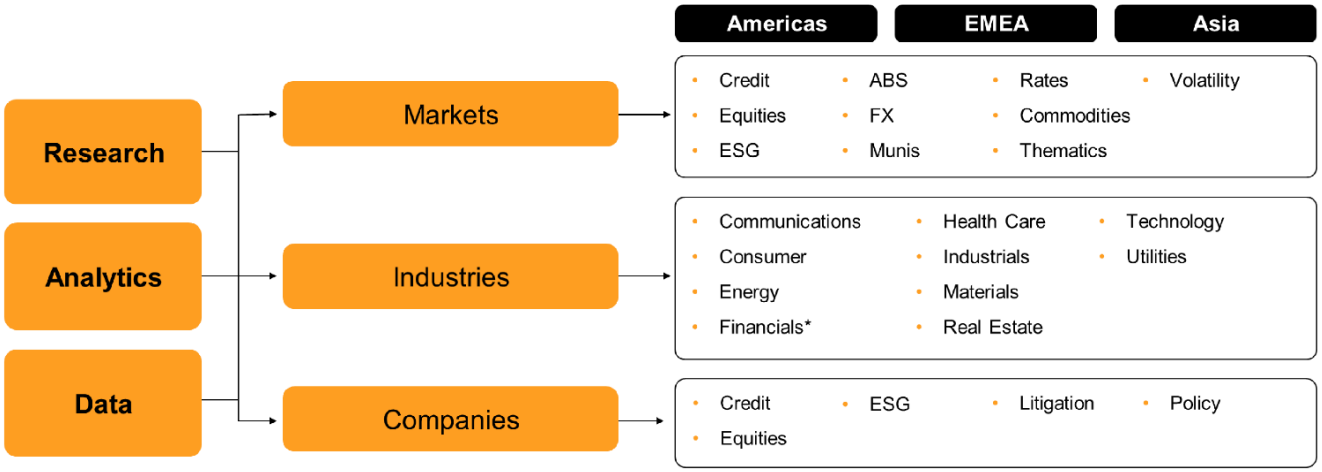
**Mumbai**
+91 22 6120 3600

**Sao Paulo**
+55 11 2395 9000

**Tokyo**
+81 3 4565 8900

# Bloomberg Intelligence

Research, analytics and data tools to help you make informed investment decisions






No portion of this document may be reproduced, scanned into an electronic system, distributed, publicly displayed or used as the basis of derivative works without the prior written consent of Bloomberg Finance L.P. Copyright and Disclaimer notice on page 98 applies throughout. © Bloomberg Finance L.P. 2025

# Bloomberg Intelligence by the Numbers.

**500**
research professionals

**135+**
industries

**600+**
data contributors

**2,000+**
companies

**21**
markets covered

