



2024

ANNUAL REPORT

AMD 
together we advance_

Dear Shareholders,

Every day, more than 27,000 AMDers around the world come to work united by a single purpose: push the limits of high-performance and adaptive computing to help solve the world's most important challenges. The impact of our work has never been more profound, and the opportunities ahead have never been greater.

This is an incredibly exciting time for AMD and the entire semiconductor industry as chips now power nearly every aspect of modern life. AMD technology is everywhere, enabling the essential devices and services used by billions of people every day.

Today, no technology better exemplifies the power of high-performance computing than AI. As the most transformative technology of the last 50 years, AI will fundamentally reshape industries and daily life in ways that extend far beyond anything we can imagine.

In the data center alone, we expect the AI accelerator total addressable market to grow by more than 60% annually, reaching \$500 billion by 2028 – equivalent to the entire semiconductor industry's annual revenue in 2023. Beyond the data center, the rapid expansion of AI across PCs, embedded devices and the edge will drive demand for more compute across each of our target markets.

Delivering on the promise of AI will require a diverse set of energy-efficient and high-performance compute engines, which we have spent years developing. However, truly addressing the scope and complexity of AI demands operational excellence at scale and a complete ecosystem built on trusted partnerships

with industry leaders. Our decade-long transformation has prepared us in each of these dimensions. We believe AMD is one of the few companies with the technical expertise and deep customer partnerships to deliver seamless end-to-end AI solutions – from the cloud to the edge and across thousands of end devices.

Our differentiated AI strategy builds on close collaboration with hardware and software leaders across the industry. We believe open standards and ecosystems accelerate innovation and broaden adoption. Just as standards-based protocols drove significant internet growth and transformed industries, establishing key standards and creating an open software ecosystem are critical to enabling the full potential of AI.

2024 marked a major milestone in our journey as we:

- Achieved record server processor market share, driven by growing adoption of AMD EPYC™ processors across cloud, enterprise and supercomputing customers.
- Scaled our data center AI business to more than \$5 billion in annual revenue, with Meta, Microsoft, Oracle and others deploying AMD Instinct™ MI300 accelerators in large-scale production environments.


- Significantly enhanced our ROCm™ AI software stack by integrating new features, optimizing popular libraries, enabling new frameworks and expanding third-party ecosystem support that increased inferencing performance of AMD Instinct MI300X accelerators 2.7x since launch.
- Gained client processor market share based on cutting-edge performance, efficiency and AI capabilities of our AMD Ryzen™ processors.
- Solidified and extended our position as the industry's leading provider of adaptive computing solutions.


Based on this momentum, we delivered strong financial results in 2024. Annual revenue grew 14% to \$25.8 billion as record Data Center and Client segment revenue more than offset declines in our Gaming and Embedded segments. Importantly, combined sales of our AMD EPYC and AMD Instinct data center products nearly doubled year-over-year, contributing approximately 50% of annual revenue. As a result, we achieved record annual revenue while significantly expanding gross margin and increasing profitability year-over-year.

2024 Financial Performance



Data Center Transformation





DATA CENTER

*Data Center segment
annual revenue increased
94% year-over-year to
a record \$12.6 billion.*

2024 was another major inflection point for our server business as share gains accelerated. EPYC processors have become the CPU of choice for the modern data center as we have consistently executed our multi-year roadmaps and delivered a top-to-bottom portfolio that offers significant performance and TCO advantages across virtually every enterprise and cloud workload.

The world's largest cloud providers have deployed EPYC CPUs at scale to power their most important and popular

services, including Office 365, Facebook, Microsoft Teams, Salesforce, SAP, Zoom, Uber, Netflix and many more. Meta alone has deployed more than 1.5 million EPYC processors across their global data center fleet.

More than 1,000 AMD-powered public cloud instances are now available from AWS, Alibaba, Google Cloud, Microsoft Azure, IBM Cloud, Oracle, Tencent and others, an increase of 27% year-over-year. Enterprise adoption of EPYC in the cloud grew significantly in the year with large businesses activating nearly 3x more EPYC cloud instances year-over-year.

Enterprise adoption of EPYC processors also surged as there are now more than 450 AMD EPYC server platforms available from the largest OEMs and ODMs. We expanded our footprint with

Forbes 2000 companies, highlighted by large wins with Adobe, Micron, Nestlé, ServiceNow, Tata, Verizon, Visa and others who deployed EPYC servers to power their mission-critical workloads.

In October, we launched our 5th Gen EPYC processor family delivering breakthrough performance and leadership TCO across enterprise scale-up and cloud-native scale-out workloads. Our latest EPYC processor family has set more than 540 performance world records across business applications, data management, virtualized infrastructure, supercomputing and digital services workloads. All of our cloud partners are planning to deploy 5th Gen EPYC processors across their infrastructure, and server providers are on-track to launch more than 120 platforms supporting the new processors.


2024 was also an outstanding year for our data center AI business.

We launched our AMD Instinct MI325X GPUs with industry-leading inferencing performance, accelerated our AI hardware roadmap to deliver an annual cadence of new accelerators, significantly enhanced our ROCm software suite and strengthened our partnerships with many of the world's largest AI providers.


Leading cloud providers and AI innovators adopted AMD Instinct MI300 accelerators at scale to power their AI workloads. Microsoft highlighted the TCO advantages of AMD Instinct MI300X accelerators when running multiple Microsoft Copilot services that use the latest GPT-4 models. Meta announced broad deployment of MI300X for inference, including serving all live traffic for its most demanding Llama 405B frontier model exclusively on AMD GPUs.

Microsoft, Oracle Cloud, IBM Cloud and more than a dozen other AI-specialized cloud providers adopted AMD Instinct accelerators to power their public cloud instances, including flagship instances available on Microsoft Azure that scale up to thousands of GPUs for AI inference and training and high-performance computing workloads.

Annual Cadence of New AI Accelerators



The El Capitan supercomputer at Lawrence Livermore National Laboratory, powered by AMD Instinct MI300A APUs, debuted as the fastest computer in the world on the latest TOP500 list. AMD proudly powers two of the top three exascale computers in the United States, and El Capitan is one of them. Adoption of our AMD EPYC and AMD Instinct processors with the broader supercomputer customer base also grew in the year, with AMD-based systems comprising 50% of the top ten fastest and 40% of the ten most energy-efficient supercomputers in the world.



We expanded functionality at every layer of our ROCm stack.

Today, more than 1 million models on Hugging Face run seamlessly on AMD hardware, and our platforms are supported natively in leading frameworks like PyTorch and JAX, emerging compilers such as OpenAI Triton, and serving solutions like vLLM and SGLang.

While we have made significant progress with our hardware and platforms, accelerating our software roadmap and strengthening broad ecosystem support remain a top priority. Our strategy is to establish ROCm as the industry's leading open software stack for AI, providing developers with greater flexibility and accelerating innovation. To support the growing adoption of AMD Instinct accelerators, we are increasing investments to enhance the out-of-the-box experience and expanding our resources supporting the open-source community, enabling faster development, testing and release of software enhancements.

In August, we acquired Silo AI, a private AI lab in Finland with extensive experience developing tailored AI solutions for Allianz, Ericsson, Finnair, Korber, Nokia, Phillips, T-Mobile, Unilever

and others. The Silo team brings deep expertise in large language model development that will enable customers to build optimized AI solutions for AMD hardware.

We took several major steps in our long-term strategy to deliver AI solutions at data center scale. To address the need for scale-up and scale-out networking performance for AI data centers, AMD is a founding member of the Ultra Accelerator Link (UALink) and the Ultra Ethernet Consortium (UEC), both of which are defining open industry standards for AI networking to allow an open ecosystem of AI accelerators to work seamlessly together.

We also announced our intention to acquire ZT Systems, a leading provider of AI infrastructure to many of the largest hyperscale computing companies. We expect ZT's deep expertise across system design, validation, networking and testing to significantly strengthen our systems and customer enablement capabilities and greatly accelerate the time to validate and deploy AMD AI solutions at scale.