Semiconductor Markets: Key Applications and Players

September 2025

SELECTED PLAYERS

- **Fabless Companies** (chipset design companies with limited or no wafer fabrication capability)
- **Integrated Design Manufacturers** (vertically integrated semiconductor companies that design, manufacture and sell chipsets)
- **Pure-Play Foundries** (manufacture chipsets exclusively for Fabless Companies and do not market or sell chipsets into downstream markets)

Selected players: Ambarella; amazon.com (Home Grown Chip Design); AMD; Analog Devices (Analog ICs, Mixed-signal); Ampere; Apple (Home Grown Chip Design); arm (Closed Chip Design Architect Only); Broadcom; Fujitsu; Google (Home Grown Chip Design); GlobalFoundries; Huawei / HiSilicon; Infineon (Microcontrollers, Power Chips); intel; intel foundry; Lattice; Marvell; MediaTek; Micron; Microchip (Microcontrollers, Mixed-signal Chips, Memory); Microsoft (Home Grown Chip Design); Mitsubishi Electric; Nordic (Wireless, Analog ICs, Power Control); Northrop Grumman; NVIDIA (GPUs, CPUs); NXP (Analog ICs, Mixed-Signal, Power Control); onsemi (Analog ICs: Power Control); Qorvo; Qualcomm (CPUs, RF, CDMA); Renesas (Analog: General, Mixed-signal); RISC-V (Open Chip Design Architect Only); RTX; Samsung; Samsung Foundry; SK hynix; Skyworks (ICs: Mixed Signal, Power Control); SMIC; ST (Microcontrollers, Analog ICs: General, Mixed-signal); Sony; Tesla; Texas Instruments (Microcontrollers, Analog ICs: General, Mixed-signal, Power); Toshiba; tsmc; UMC; Western Digital

KEY SEMICONDUCTOR APPLICATIONS

AUTOMOTIVE
- Advance Driver Assistance Systems (ADAS)
- Autonomous Vehicles
- Digital Cockpit (in-vehicle infotainment)
- Vehicle Motion and Dynamics Systems

NETWORKING AND COMMUNICATIONS
- Networking Equipment and Fiber to the Home (FTTH) (connectivity solutions)
- Wireless Solutions (cellular, IoT, Bluetooth)
- Satellite Communication Systems (SATCOM)
- Radio Frequency Integrated Circuits (transceivers and radio frequency front end)
- 5G Telecom Networks

INDUSTRIAL, IoT AND ROBOTICS
Access Control, Warehouse Management and Logistics, Smart Sensors, Industrial Machine Vision, Occupancy Monitoring, and Elder Care
- Energy (Crude Oil, Natural Gas, NGL, and Refined Petroleum Products Pipelines)
- Robots and Robotic Arms
- Factory Automation Systems
- Building Control (optical devices/smart vision, security)
- Power Storage, Generation & Transmission (wind power, HVDC, solar cells)
- Telemedicine and Medical Devices (e.g., industrial diagnostic equipment)

CONSUMER
Appliances, Consumer Electronics, HVAC, and Smart Home IoT
- Appliances (washing machines, refrigerators, microwaves, dishwashers, consumer robots)
- TVs, Speakers, Displays and Projectors
- Heating, Ventilation and Air Conditioning (HVAC)

RETAIL
- Commercial Self-Service Devices (ATMs, point-of-sale devices, biometrics, subway channel gates, mobile payments)

AEROSPACE & DEFENSE
- Military Aircraft, Drones, Rotorcraft, Missiles, and Weapons Systems (control, monitoring, communication, navigation, safety, and emergency systems)
- Commercial Aerospace (control, monitoring, communication, navigation, safety, and emergency systems)

INFORMATION TECHNOLOGY
Hyperscale Data Centers/Cloud Computing, Core Computing, and Edge Computing
- Servers: AI Computing (for the AI data center)
- Servers: High-Performance Computing
- Servers: Standard Computing
- Servers: Edge Computing
- Data Storage (compression, duplication, data reconstruction)

Personal Computers
- Desktops and Laptops (includes gaming desktops and laptops)
- Tablets
- Smart Phones
- Videogame Consoles and Mobile Gaming Devices
- Extended Reality Headsets (AR, VR, MR, XR)
- Wearables (smart watches and fitness trackers)

Foundry and chip design notes

arm → Foundry Manufacturing and Assembly Services (customers include Ambarella, Renesas, Cadence, NetFoundry, NVIDIA, Marvell, Qualcomm, Siemens, Microsoft, Fujitsu, Google, Ansys, Toshiba, AMD, Apple, NXP, ST)

Arm Holdings plc is the leading technology provider of processor IP, offering the widest range of processors to address the high-performance, power/energy-efficient, and low-cost requirements of every device. Its primary business is the design of central processing unit (CPU) cores that implement the ARM architecture family of instruction sets. It also designs other chips, provides software development tools and provides systems IP, physical IP and platforms, system-on-chip (SoC) infrastructure, and software.

Over 70% of the world's population uses ARM-based products. ARM processor designs run most of the world's software, such as operating systems and applications for smartphones, tablets and personal computers, data centers and networking equipment, vehicles, and embedded operating systems in devices such as smartwatches, thermostats, drones and industrial robotics.

ARM's customers are many of the world's leading semiconductor companies (e.g., Advanced Micro Devices, NVIDIA, Qualcomm, and Samsung Electronics) and original equipment manufacturers (e.g., Amazon, Apple, Google, and Microsoft). These customers put ARM's processor designs in physical chips that will be incorporated into a wide variety of computing devices. ARM is majority owned by Japanese conglomerate SoftBank Group.

GlobalFoundries → Foundry Manufacturing and Assembly Services (AMD, Infineon, MediaTek, NXP, Qualcomm)

Global Foundries is the world's fourth largest chipmaker (by revenue). However, Global Foundries does not focus on leading-edge chips (which are the focus of Taiwan Semiconductor Manufacturing Company), such as powerful graphics processing units used to train large language models like ChatGPT. Instead, the company manufactures chips that perform functions like power management, connecting to displays, or enabling wireless connections.

intel foundry → Foundry Manufacturing and Assembly Services (Broadcom, Ericsson, intel, MediaTek, Microsoft)

Intel Foundry is a newly established operating segment of Intel that includes foundry technology development, foundry manufacturing and supply chain, and foundry services. Intel's seeks to become a world-class foundry with the goals of being the number two foundry by 2030, delivering five new process nodes by 2030, and expanding its lead in advanced processing.

In addition to its own manufacturing facilities, Intel Foundry is partnering with (1) Tower Semiconductor, which will be investing in new equipment to be installed in Intel's New Mexico facility for analog products, and (2) UMC, which will leverage three of Intel's older Arizona fabs and process nodes to support applications like industrial IoT, mobile, communications, infrastructure, and networking.

RISC-V → (Alibaba, Quintauris, Microchip)

Semiconductor industry players Robert Bosch, Infineon Technologies, Nordic Semiconductor, NXP Semiconductors, and Qualcomm established Munich, Germany-based *Quintauris GmbH* to advance the adoption of RISC-V globally by enabling next-generation hardware development, with the aim to be (1) a single source to enable compatible RISC-V-based products, provide reference architectures, and help establish solutions to be widely used across various industries. The initial application focus will be automotive, but with an eventual expansion to include mobile and IoT.

RISC-V International is a non-profit organization that promotes the adoption and development of an open standard instruction set architecture (ISA) called Reduced Instruction Set Computer – V (or RISC-V). Founded to provide a free and open ISA that is extensible and suitable for various computing applications, RISC-V International oversees the development of the RISC-V ecosystem, including hardware implementations, software tools, and community engagement.

RISC-V International aims to drive innovation and collaboration among hardware and software developers, ensuring that RISC-V remains a competitive and accessible alternative to proprietary ISAs, such as ARM plc's proprietary and tightly-controlled ISA) that offers established reliability and compatibility but limits customization.

By contrast, open ISAs, exemplified by RISC-V, are community-driven and provide greater flexibility for customization, fostering innovation and adaptation to specific needs. RISC-V is well-suited for embedded systems, such as microcontrollers, IoT devices, and industrial control systems, due to its simplicity, efficiency, and customizable instruction set.

Samsung Foundry → Foundry Manufacturing and Assembly Services (Ambarella, AMD, Meta, Google, IBM, intel, NVIDIA, NXP, Qualcomm, Samsung, Sony, ST, Tesla)

Founded in 2005, Samsung Foundry is the world's second largest semiconductor foundry company. The company offers a full suite of competitive, advanced process technologies, design technologies, IP, and high-volume manufacturing capability to over 150 customers through its facilities in South Korea and Texas.

Samsung Foundry also offers advanced packaging solutions and application-specific technologies for radio frequency or RF, eNVM, high voltage, BCD, CMOS image sensors, and fingerprint sensors. Along with TSMC, Samsung is able to produce 3-nanometer chips; the smaller the nanometer size, the more advanced and efficient the chip.

SMIC → Foundry Manufacturing and Assembly Services (Broadcom, Huawei/HiSilicon, Qualcomm, Texas Instruments)

State-backed Semiconductor Manufacturing International Co. (SMIC) is China's largest chipmaker and the world's third-largest foundry (by revenue). The company's chips are found in automobiles, smartphones, computers, and IoT technologies. China consumes nearly 50% of the world's semiconductors because it is the biggest assembly market for consumer devices. Huawei/Hisilicon is SMIC's largest customer, accounting for 19% of 2023 revenues.

In light of the U.S. sanctions that restrict the ability of U.S. chipmakers to sell certain U.S. technology and chips in China, SMIC is critical to China's hopes of reducing dependence on foreign technology in its domestic semiconductor industry. As a result, the Chinese government has invested billions in yuan in subsidies in Chinese chip makers (including SMIC) to boost domestic chip production.

tsmc → Foundry Manufacturing and Assembly Services (Analog Devices, Broadcom, ST, AMD, Apple, MediaTek, Marvell, Qualcomm, NVIDIA, Sony, Tesla)

Taiwan Semiconductor Manufacturing Company (TSMC) is the world's second-most valuable semiconductor company, the world's largest dedicated independent ("pure-play") semiconductor foundry, and its country's largest company, with headquarters and main operations located in the Hsinchu Science Park in Hsinchu, Taiwan.

Chips for PCs and smartphones account for most of TSMC's revenue. TSMC also focuses on manufacturing sophisticated GPUs essential to training AI models. Patents around TSMC's chip-packaging process give TSMC a competitive advantage. There are plenty of companies in the semiconductor foundry space, but none of the competitors' chips compares to the advanced capabilities of TSMC chips.

TSMC's competitive advantage lies in the end stages of chip production and a process called "packaging"—i.e., the process of enclosing the chip and putting in electrical connections vital for the chip's ability to function in different electronics. The packaging process is done by chip producers, but TSMC has an extensive list of patents (approximately 3,000) that protect its advanced chip-packaging technology, which enables TSMC to essentially stack and combine multiple chip components into a single item, leading to higher performance than others.

UMC → Foundry Manufacturing and Assembly Services (AMD Xilinx, Broadcom, Infineon, MediaTek, Qualcomm, Realtek, Texas Instruments)

United Microelectronics Corporation (UMC) designs, manufactures, and markets integrated circuits and related electronic products. UMC was founded in 1980 as Taiwan's first semiconductor company as a spin-off of the government-sponsored Industrial Technology Research Institute (ITRI).

United Microelectronics UMC is headquartered in Hsinchu, Taiwan also has facilities and local offices in the U.S., Europe, China, Japan, South Korea, and Singapore. UMC has total of 12 fabs in production with combined capacity of more than 400,000 wafers per month. The company has a worldwide total of 20,000 employees.

As a pure-play foundry, UMC provides high quality integrated circuit (IC) fabrication services, focusing on analog chips (such as logic, mixed-signal, and embedded high voltage chips), memory chips and various specialty technologies to serve all major sectors of the consumer electronics, displays, personal computer peripherals, automotive, and communications industries.

AMD: Overview

September 2025

DATA CENTER

(FY2024 Net Revenue: $12.6B)

This segment designs and sells the specialized processors that perform the intensive computational tasks required by data centers, primarily AI accelerators, server CPUs, GPUs, APUs, DPUs, FPGAs, SmartNICs, and Adaptive SoC products for data centers

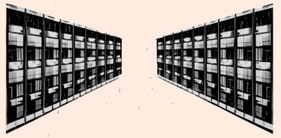

Server Central Processing Units (CPUs)	Field Programmable Gate Arrays (FPGA) & Adaptive System-on-Chip (SoC) Data Center Products	Networking Products	Data Center Graphics Processing Units (GPUs)
AMD EPYC	FPGA: AMD ARTIX 7, AMD KINTEX, AMD VIRTEX, AMD SPARTAN Adaptive SoCs: AMD ZYNQ, AMD ZYNQ MPSoC, AMD VERSAL, AMD Vivado	AMD PENSANDO AMD SOLARFLARE	AMD INSTINCT

CLIENT

(FY2024 Net Revenue: $7.1B)

This segment primarily includes (1) central processing units (CPUs) from AMD's Zen line, (2) accelerated processing units (APUs) that integrate CPUs and GPUs, and (3) chipsets for desktop and notebook personal computers.

Desktop CPUs and APUs	Notebook CPUs	Commercial CPUs	Chipsets (for Desktop PCs)
AMD RYZEN AMD ATHLON ACCELERATED A-SERIES AMD	AMD RYZEN AMD ATHLON	AMD RYZEN AMD RYZEN THREADRIPPER	WRX80 AMD SOCKET AM5 X399 TRX40 AMD 9-SERIES CHIPSET AMD A SERIES

GAMING

(FY2024 Net Revenue: $2.6B)

This segment Primarily includes graphics solutions to enable immersive visualization and to process AI/ML based workloads for use in data centers, desktop PCs, notebook PCs, handheld devices, All-in-Ones, and professional workstations along with semi-custom solutions.

Semi-Custom SoC Products	Discrete Desktop and Notebook GPUs	Professional GPUs (for integration in mobile and desktop workstations)
XBOX STEAM DECK	AMD RADEON (GPU processors that support application programming interfaces and high refresh rate displays, software, and graphics cards)	AMD RADEON PRO

EMBEDDED

(FY2024 Net Revenue: $3.6B)

This segment provides specialized computing solutions such as embedded CPUs, GPUs, APUs, FPGAs, and Adaptive SoC products for devices that aren't traditional computers in applications such as industrial and automation, medical, automotive, and networking and communications.

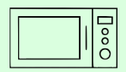

Embedded CPUs, APUs, and GPUs	FPGAs and Adaptive SoCs	Development Boards and Kits and Adaptive SoC Products	Legacy Product Families
AMD EPYC EMBEDDED AMD RADEON AMD RYZEN Embedded	FPGA Families: AMD ARTIX 7, AMD KINTEX, AMD SPARTAN Adaptive SoC Products: AMD ZYNQ, AMD VERSAL, AMD ZYNQ MPSoC	AMD KRIA Configuration Products: AMD Vitis, AMD Vivado	AMD VIRTEX AMD SPARTAN CoolRunner-II XC9500XL

AMD: Ecosystem

Advanced Micro Devices, Inc.
2485 Augustine Drive
Santa Clara, California 95054
Phone: (408) 749-4000 ; Website: www.amd.com

Outside Relationships

Regulators

General Business Regulation

Environment

U.S. Environmental Protection Agency (EPA) (regulation of, and permits for, discharges, leaks, emissions, disposals, storage of particular matter, hazardous waste, and toxic substances in the water, air and soil; superfund cleanup requirement for contamination)

Labor and Employment

U.S. Department of Labor (DOL) (***DOL*** regulation of work practices, wage and overtime record-keeping; ***Equal Opportunity Employment Commission*** enforcement of anti-discriminatory laws in hiring, workplace practices and civil rights; ***Occupational Safety and Health Administration*** workplace health and safety, and serious injury reporting)

Intellectual Property

U.S. Patent and Trademark Office (USPTO); U.S. Copyright Office (USCO) (patent, trademark, service mark, copyright registrations)

Restricted Transactions

U.S. Department of the Treasury (***FinCen*** "Know Your Customer" data collection requirements / anti-money laundering regulation; ***OFAC*** restricted party transactions regulation; and ***Treasury*** economic sanctions)

Federal Taxes

U.S. Internal Revenue Service (IRS) (regulation of corporate income, withholding and payroll taxes)

Antitrust and Fair Trade

U.S. Department of Justice (DOJ) and Federal Trade Commission (FTC) (***DOJ/FTC*** enforcement of antitrust and fair-trade laws, and business combination/merger reviews and clearances; ***FTC*** regulation regarding data privacy; and ***DOJ*** Foreign Corrupt Practices Law enforcement)

Imports

U.S. Department of Homeland Security / Customs and Border Protection (CPB) (regulation of imports of finished goods, raw materials and supplies; customs duties administration)

Exports

U.S. Dept of Commerce Bureau of Industry and Security / Export Administration (regulation of exports of finished goods, raw materials and supplies, and information handling; export licenses)

Capital

Public Debt Holders

Hedge Counterparties

Commercial Banks (Administrative Agent: Wells Fargo Bank)

Other Banks: Bank of America, JPMorgan Chase Bank, Credit Suisse, Goldman Sachs Bank, Morgan Stanley Bank, MUFG Bank, Bank of China, Citibank, Bank of New York Mellon, and Standard Chartered Bank)

Suppliers

Key AMD Primary Manufacturing Suppliers

Assembly, Test, Marking, and Packaging (ATMP) Services

ATMP Joint Ventures with Affiliates of Tongfu Microelectronics Co., Ltd. (AMD owns a 15% equity interest)

(wafers for AMD's products are either sorted by the foundry or delivered by the foundry to its ATMP partners or subcontractors located in the Asia-Pacific region who package and test AMD's final semiconductor products; AMD's FY2024 purchases from the JV: $1.7B)

Memory and Other Technologies

(produces AMD's integrated circuits in the form of programmable logic devices)

(produces (1) wafers for AMD's HPC, FPGA and Adaptive SoC products and (2) integrated circuits in the form of programmable logic devices)

Customers

Data Center Customers

Hyperscale Data Centers (to operate web-based applications or to support public cloud computing and storage service offerings)

Original Equipment Manufacturers (OEMs) (manufacturers of servers, storage arrays, workstations, networking equipment, desktops, notebooks, PC motherboards and game consoles)

Original Design Manufacturers (ODMs) DESIGN ENGINEERING SERVICE PROVIDER

System Integrators

Independent Distributors

Add-in-Board Manufacturers (AIBs)

Business Regulation
Bond Financing
Derivatives to Mitigate Foreign Currency Exchange Risk
Working Capital Financing

Sales of AI Accelerators, Server CPUs, GPUs, APUs, DPUs, FPGAs, SmartNICs, and Adaptive SoC Products

Supplies of (1) Assembly, Test, Marking and Packaging Services, (2), Memory and Other Technologies, (3) Foundry Services for Wafers/Chipsets and Integrated Circuits, and (4) Intellectual Property (via third-party licenses)

Advanced Micro Devices, Inc. (a Delaware corporation)

Debt Structure

Long Term Debt ($1.7B @ 12/28/2024) | Credit Ratings: A2 (Moody's), A- (S&P)

Working Capital

Revolving Credit Facility: $3.0 B Capacity, Outstanding: $0 (Matures 2027)

Commercial Paper Program: $3.0 B Capacity, Outstanding: $0

Debt Issued or Guaranteed by Parent

2030 Senior Notes: $750M @ 2.375%

2032 Senior Notes: $500M @ 3.924%

2052 Senior Notes: $500 M @ 4.393%

Logo History

Corporate Brands

1969-1995 | 1980 – 1990 (alternate logo) | 1995 – Present

Equity Structure

Equity Capital

Share Repurchase Program
Authorized: $12.0B
Expiration: None
Balance: $4.7B

Common Stock
Authorized: 2.2B shares
Issued: 1.7B shares
Outstanding: 1.6B shares
Record Holders: 5,028

Governance

Board of Directors

- Dr. Lisa T. Su
- Nora M. Denzel (I, N (chair))
- Mike P. Gregoire (C (chair), I)
- Joseph A. Householder (A (chair), N)
- John W. Marren (A)
- Jon A. Olson (A)
- Abhi Y. Talwalkar (C, I (chair))
- Elizabeth W. Vanderslice (C, N)

Committees
Audit and Finance (A), Compensation and Leadership Resources (C), Innovation and Technology (I), Nominating and Corporate Governance (N)

Executive Leadership

Chair and CEO: Dr. Lisa T. Su	EVP and GM, Data Center Solutions: Forrest Norrod	SVP, Chief Human Resources Officer: Robert Gama
EVP, Chief Sales Officer and President AMD EMEA: Darren Grasby	EVP, Chief Technology Officer: Mark Papermaster	SVP, General Counsel and Corporate Secretary: Ava Hahn
EVP, Chief Commercial Officer: Phil Guido	SVP, Chief Administrative Officer: Ruth Cotter	SVP and Chief Strategy Officer, Corporate Development: Mathew Hein
EVP, CFO and Treasurer: Jean Hu	SVP, Global Operations: Keivan Keshvari	

Corporate Matters

Human Resources
- Talent Acquisition, Development and Management
- Total Rewards
- Mission, Culture, and Engagement
- Global HR Team
- Diversity, Belonging, and Inclusion

Sales
- Worldwide Sales Organization (end-to-end product portfolio, spanning data center, commercial, consumer and embedded markets)

Finance and Accounting
- Accounting
- Global Corporate Services
- Treasury
- Financial Reporting
- Tax
- Global Finance Organization
- Risk Management
- Indirect Procurement
- Investor Relations
- Facilities
- Strategy
- Internal Audit

Technology

Technology
- Technical Direction
- Engineering
- Product Development (Microprocessor Design, I/O and Memory, System-on-Chip or SOC Methodology)

IT
- Compute Infrastructure
- IT Services

Business Operations
- End-to-End Manufacturing Strategy
- Supply Chain Strategy

Legal
- Global Legal Operations
- Corporate Governance
- Litigation Management
- Intellectual Property
- Ethics/Compliance Programs
- Public Affairs (government relations, community affairs, and corporate responsibility)

Strategy
- Strategic Planning
- Execution Initiatives
- Venture Investments
- Acquisitions

Marketing, HR, and Investor Relations
- Transaction Integration Planning Office
- HR Policy and Planning
- Global Marketing Team
- Brand Awareness and Management
- Advertising
- Investor Relations

Client Business Unit
- Client Computing Products and Solutions (product strategy, management, and engineering)

Operations
- Supply Chain Strategy and Operations
- Capacity Planning
- Major Category Procurement
- Global Technology Business Operations
- New Product Introduction Engineering
- End-to-end Manufacturing

Administration
- Marketing
- Communications
- Human Resources
- Corporate Security
- Information Technology
- Indirect Procurement

Operations

Data Center (FY2024 Net Revenue: $12.6B)

The Data Center segment designs and sells specialized processors that perform the intensive computational tasks required by data centers. These processors include primarily AI accelerators, server CPUs, GPUs, APUs, DPUs, FPGAs, SmartNICs, and Adaptive SoC products for data centers.

Product Brands

Server CPUs
Server CPUs are the workhorses of the data center, engineered to run continuously, handle thousands of tasks simultaneously, and manage vast amounts of data for many concurrent users.

Higher Core Counts: They have far more processing cores to handle multiple parallel workloads.
Larger Memory Capacity: They support more system memory (RAM) and can use special Error-Correcting Code memory to prevent data corruption.
Enhanced Reliability: Built for 24/7 operation and have specialized features to ensure stability and uptime.
Multi-Socket Support: Can be used in configurations with multiple CPUs on a single motherboard.

FPGAs and Adaptive SoCs (system-on-chip)
Customizable and flexible hardware

FPGAs: Chips that can be rewired or reprogrammed after being manufactured. They're not a CPU or a GPU, but a blank slate of programmable logic that can be configured to execute specific tasks.

Adaptive SoCs: Combines a traditional processor core with a programmable FPGA fabric on a single chip. Instead of relying on a general-purpose CPU, a user can program an FPGA to perform a specific function—like video encoding—with a level of speed and efficiency that a CPU or even a GPU can't match.

Networking Products
Includes DPUs and SmartNICs, specialized processors that offload networking, security, and storage tasks from the main CPUs.

Pensando DPUs: Along with comprehensive networking software stack offload data center infrastructure services from the host CPU, used by large infrastructure as a Service (IaaS) cloud providers to accelerate workload performance for hosted virtualized and bare-metal offerings.

AMD Solarflare Portfolio: A low latency networking solution combining hardware adapters with the Onload family of user space networking libraries.

Data Center GPUs

While GPUs were originally designed for graphics, data center GPUs are specialized accelerators that handle the heavy lifting for artificial intelligence, machine learning, big data analytics, and high-performance computing. AMD Instinct accelerators perform these tasks at a fraction of the time and power cost of CPUs, making them the key to modern AI and scientific research.

Client and Gaming ($9.6B; Beginning in FY2025, AMD changed its segment structure, combining the Client and Gaming segments into one reportable segment)

Client ($7.1B)

The Client business unit primarily includes (1) central processing units (CPUs) from AMD's Zen line, (2) accelerated processing units (APUs) that integrate CPUs and GPUs, and (3) chipsets for desktop and notebook personal computers.

Products

Desktop CPUs and APUs
High Performance: Optimized for speed and raw computing power and typically have a higher power consumption (and generate more heat); designed to be used in a larger case with robust cooling systems.
Physical Size: Larger than their mobile counterparts, built to fit into a standardized socket on a motherboard.
User Upgrades: Generally easy for a user to install, allowing PC builders to customize their systems for specific tasks like gaming or video editing.

Notebook CPUs
Power Efficiency: Its top priority is to consume as little power as possible to extend battery life. They are designed to run cooler and can adjust their performance based on the workload to save energy.
Integrated Design: They are often soldered directly onto the motherboard, making them much more compact and difficult or impossible to upgrade, allowing for thinner and lighter laptop designs.
Integrated Graphics: Many notebook CPUs include an integrated graphics processing unit on the same chip, eliminating the need for a separate graphics card in many laptops. This also saves space, power, and cost.

Commercial CPUs
Similar to desktop/notebook CPUs with additional security and management features for the corporate market.

Security Features: Often include specialized hardware and software for enhanced security, remote management, and data protection.
Reliability and Stability: Tested to ensure they can run reliably in a business environment for a longer product lifecycle and software stability.
Remote Management: Often come with features that allow IT professionals to remotely manage, update, and troubleshoot computers at remote locations.

Chipsets (for Desktop PCs)
A collection of integrated circuits on the motherboard whose job is to manage the flow of data between the CPU and all the other components.

Communication Hub: Connects the CPU to everything else: the RAM (memory), the GPU (graphics card), storage drives (SSDs and HDDs), and external devices through USB ports, Wi-Fi, and other connections.
Feature Control: Determines which features the motherboard supports. For example, it dictates how many USB ports are available.
Compatibility: Must be compatible with a specific generation of CPUs; AMD Ryzen CPU on an Intel motherboard.

Gaming ($2.6B)

The Gaming business unit primarily includes graphics solutions to enable immersive visualization and process AI/ML based workloads for use in data centers, desktop PCs, notebook PCs, handheld devices, All-in-Ones, and professional workstations along with semi-custom solutions.

Products

Semi-Custom Products
Tailored, high-performance, customer-specific solutions based on AMD CPU, GPU and multi-media technologies

Semi-Custom SoC Products
AMD works directly with clients to design a tailor-made, all-in-one chip for their unique product. AMD products power the Sony PlayStation 5, Microsoft Xbox Series S and X game consoles, and the Valve Steam Deck PC.

While AMD provides the underlying technology, the gaming companies themselves typically don't use a publicly known AMD brand name like "Ryzen" or "Radeon" for the final product. Instead, the chips are often referred to by the console's branding, as they are a custom, collaborative creation.

Discrete Desktop and Notebook GPUs
"Discrete" means the GPU is a separate, dedicated component from the CPU. They are sold both as individual graphics cards for users to build their own PCs and as components within pre-built gaming desktops & laptops.

AMD Radeon Series Discrete GPU Processors for Desktop and Notebook PCs: support current generation application programming interfaces like DirectX 12 Ultimate and Vulkan, support high-refresh rate displays and include the latest technologies for immersive gaming experiences and high-performance AI/ML computing.

Professional GPUs
The Professional GPUs business unit consists of "discrete" GPUs, but they are built for professionals in digital content creation, engineering, and architecture industries.

AMD Radeon PRO Family of Professional Graphics Products: Designed for integration in desktop workstations optimized through hardware and software for demanding use cases such as 3D rendering, design and manufacturing for Computer-Aided Design, and media and entertainment for broadcast and animation pipelines on high resolution displays.

Embedded ($3.6B)

The Embedded Segment provides specialized computing solutions (e.g., embedded CPUs, GPUs, APUs, FPGAs, and Adaptive SoC products for devices that aren't traditional computers) for applications such as industrial and automation, medical, automotive, and networking and communications.

Product Brands

Embedded CPUs, APUs and GPUs

Embedded Processors and GPUs
These devices are designed to support high performance compute, high-bandwidth network connectivity and security, high-performance storage requirements for enterprise and cloud infrastructure, 3D graphics performance, and 4K multimedia requirements of automotive infotainment systems. They encompass core processors for embedded systems, but they are very different from their desktop or data center counterparts. They are designed for reliability, low power consumption, and long-term availability (often 10+ years), which is critical for industrial and medical equipment.

FPGAs and Adaptive SoCs
Customizable devices that can be tailored to meet the specific needs of each customer

FPGA Families
FPGAs are unique because their hardware can be reprogrammed after they've been manufactured, allowing for custom logic circuits.

Adaptive SoC Products
Adaptive SoCs combine a processor core (like an Arm-based CPU) with a programmable FPGA fabric on a single chip. They are used in mission-critical applications where high performance, low latency, and custom hardware acceleration are a must, such as in telecommunications, aerospace, and defense.

Development Boards, Kits and Configuration Products
This category isn't sold to end users, but rather to the engineers and developers who are building products. These tools enable them to work with AMD's embedded hardware.

Development Boards and Kits for all AMD FPGA and Adaptive SoCs
These are pre-built circuit boards with an AMD chip already installed. They allow developers to quickly prototype and test their designs without having to build their own hardware from scratch.

Configuration Products
This includes the software and tools that engineers use to program and configure AMD FPGAs and Adaptive SoCs.

Legacy Product Families
These products are no longer actively developed, but AMD provides long-term support, drivers, and documentation for them. This ensures that customers with long product lifecycles can continue to get support for their hardware for years after it was released.

Prior generation high-end Virtex FPGA families and the original Virtex family

Prior generation low-end Spartan FPGA families and the original Spartan families

Prior generations of Complex Programmable Logic Devices include the CoolRunner-II and XC9500 product families

Data Center Competition

CPU, GPU and DPU server products and FPGA and Adaptive SoC server products | CPU, GPU and DPU server products only

Client Competition

Gaming Competition

Embedded Competition

Outside Relationships

Customers

Licensees of AMD Intellectual Property

Embedded Customers
- Aerospace and Defense
- Automotive
- Industrial
- Vision and Healthcare
- Communication Infrastructure
- Audio, Video and Broadcasting
- Emulation and Prototyping (3d software for simulation and emulation)
- Test Measurement
- Data Center (includes networking, security and storage)

Sales of Embedded CPUs, GPUs, FPGA and Adaptive SoC Products, Development Boards, Kits, and Configuration Products, and Legacy Product Families

Wholesale Sales of Embedded CPUs, GPUs, FPGA and Adaptive SoC Products, Development Boards, Kits, and Configuration Products, and Legacy Product Families

Sales of Embedded CPUs, GPUs, FPGA and Adaptive SoC Products

Distributors and OEM Partners

Suppliers

Securities Law Regulation and Stock Listing Requirements

Professional Service Firms
- Ernst & Young (audit services)
- W20 (advertising services)
- Forbes Tate Partners (lobbying services)
- Skadden, Arps, Slate, Meagher & Flom LLP (transactional legal services)
- Google Cloud (cloud computing, storage, and data analytics services platform)

Capital

Equity Capital

Dividends and Common Stock Repurchases

Professional Services

Licenses of AMD Intellectual Property

Significant Shareholders
- Vanguard Group (9.53%)
- Blackrock Inc. (8.48%)
- State Street Corporation (4.44%)
- Geode Capital Management (2.26%)
- Morgan Stanley (1.82%)
- Institutional Ownership (69.20%)

Regulators

Securities Regulators
- U.S. Securities and Exchange Commission (regulation of offers and sales securities, material event disclosure and financial reporting requirements; recordkeeping requirements under anti-bribery law)
- Nasdaq Stock Market (listing, maintenance and corporate accountability rules)
- Delaware Secretary of State (administration and enforcement of Delaware corporation law governing business entity formation and dissolution, shareholder, board and officer governance / meetings, fiduciary duties of loyalty and care of directors, executive management duties, financial reporting, charter document amendments, mergers and acquisitions, and compliance / penalties for non-compliance)

Key Company Data (as of 12/28/2024)

Business Overview: Santa Clara, California-based Advanced Micro Devices Inc. (AMD) is a U.S. multinational designer of computer semiconductors and related technologies for consumer and enterprise markets. While AMD at first fabricated its own processors, AMD spun-off control of its foundries in 2009 under the "Global Foundries" name and now outsources its semiconductor fabrication work to Global Foundries. AMD develops microprocessors, embedded processors, graphics processors, motherboard chipsets, and field-programmable gate arrays for personal computers, workstations, gaming devices, embedded systems applications, servers, and data centers. AMD's technologies advance the future of data centers, powering the cloud services that have become an essential part of how individuals work, game, and connect to network, PCs, edge computing, and artificial intelligence.

Share Data
- Stock Exchange (NASDAQ) Ticker Symbol: AMD
- Share Price: $168.58 (8/28/2025)
- Earnings Per Share (Trailing 12 Months): $1.67
- Forward Annual Dividend: n/a
- Market Capitalization: $273.6B (8/28/2025)

Resources
- Employees: 28,800
- Headquarters: Santa Clara, CA
- Significant Operations Austin, Texas; San Jose, California; Shanghai, China; Markham, Ontario, Canada; Longmont, Colorado; Dublin, Ireland; Singapore; Bengaluru, India; Hyderabad, India

Financial Highlights

Statement of Earnings
- Net Revenues: $25.8B
- Cost of Sales: $13.1B
- Gross Profit: $12.7B (49.4% Margin)
- Operating Expenses: $10.8B
- Operating Income: $1.9B
- Net Income: $1.6B

Balance Sheet
- Total Assets: $69.2B
- Total Liabilities: $11.7B
- Long-Term Debt: $1.7B
- Equity: $57.6B

Cash Flow
- From Operations: $3.0B
- Used in Investing: $1.1B
- Used in Financing: $2.1B

Financial Returns
- Return on Equity: 4.70%

Key Developments

Acquisitions. On March 31, 2025, AMD acquired ZT Systems Group Int'l, Inc. for $4.4 billion. AMD then sold ZT Systems' server manufacturing business to Sanmina for $3 billion, keeping the ZT System's AI systems design and customer enablement teams. Sanmina will handle the manufacturing of AMD's AI solutions.

Client Customers

Sales of Alveo Products (Alveo is a family of high-performance AMD accelerator cards designed to accelerate workloads in data centers, cloud computing environments or at the edge; they are specialized hardware that can be plugged into a server to significantly speed up specific, computationally intensive tasks)

Value Added Resellers (VARs) and Integrated Solution Vendors (ISVs)

Sales of CPUs, APUs, and Chipsets for Desktop and Notebook Personal Computers

Personal Computer OEMs | Independent Distributors | ODM Motherboard Manufacturers (chipset products only)

Gaming Customers

Sales of Graphics Solutions (semi-custom SOC products, discrete desktop and notebook GPUs, and professional GPUs)

Personal Computer OEMs | Independent Distributors | Add-In Board Manufacturers | Independent Game Console and Portable Gaming Devices Manufacturers | Contract Manufacturers (for AMD branded graphics cards)

GlobalFoundries Inc. (produces wafers for AMD's HNL products at the 12 nm and 14 nm technology nodes)

GlobalFoundries was a part of AMD prior to its spinoff in 2009. Now, AMD outsources its less-sophisticated semiconductor fabrication work to GlobalFoundries, which is in contrast to Intel, which is an integrated device maker and currently expanding its own foundry capacity.

Intel Intellectual Property

In the 1980s, AMD was a secondary source of chip supply for Intel. However, after AMD and Intel parted ways, AMD reverse engineered Intel's chips to make its own products that were compatible with Intel's groundbreaking x86 software. Intel sued AMD and, as part of a settlement in 1995, gave AMD the right to continue designing x86 chips, making personal computer pricing more competitive for end consumers. It is believed that some advanced features may require an x86 license from Intel.

September 2025

AMD: Supplement

September 2025

KEY PLAYERS BY INDUSTRY SECTOR

Software

EDA*

ORIGINAL EQUIPMENT MANUFACTURERS

Google, amazon.com, Lenovo, Microsoft, Hewlett Packard Enterprise, SONY, SAMSUNG, JUNIPER NETWORKS, hp, DELL Technologies, cisco, IBM, NetApp

SEMICONDUCTOR FOUNDRIES

Pure Play: tsmc, GlobalFoundries, SMIC, UMC

Integrated: Micron, RENESAS, intel, intel foundry, FUJITSU, SAMSUNG, SAMSUNG Foundry, infineon

CHIP DESIGNERS

NVIDIA, Qualcomm, RENESAS, intel, TOSHIBA, BROADCOM, SAMSUNG, FUJITSU, NXP, infineon, life.augmented

OEM Designers

aws, Microsoft, Google

PROCESSOR DESIGN, IP, SOFTWARE, AND TOOLS DEVELOPERS

RISC-V, arm, intel

Source of Concept: Arm Holdings Preliminary Prospectus on SEC Form F-1, p.4.

GLOSSARY

Accelerator Cards: Accelerator cards, also known as hardware accelerator cards or coprocessor cards, are specialized hardware components designed to offload and accelerate specific tasks or workloads in a computer system. These cards are typically used to enhance the performance of a computer or server for specific types of computations, such as graphics processing, artificial intelligence, machine learning, cryptography, and more. Accelerator cards can significantly speed up these tasks by taking on the processing load, leaving the CPU (Central Processing Unit) to handle other tasks. Common types of accelerator cards are GPUs, FPGAs, and NICs.

AI: Refers to "artificial intelligence," which is the simulation of the human brain by machines or computer systems to enable problem solving. Specific applications of AI include expert systems, natural language processing, speech recognition, and machine vision.

AMD: Refers to Advanced Micro Devices, Inc., a Delaware corporation. Headquartered in Santa Clara, California, AMD is a U.S. multinational semiconductor firm that designs computer processors and related technologies for consumer and enterprise markets. While the company at first made its own processors, AMD spun-off its foundries in 2009 under the Global Foundries name and now outsources its semiconductor fabrication work to GlobalFoundries and TSMC. The main products developed by AMD include microprocessors, embedded processors, graphics processors, motherboard chipsets, and field-programmable gate arrays (FPGAs) for personal computers, workstations, gaming devices, and embedded systems, applications servers, and data centers.

APIs: Application Programming Interface, which is a software intermediary, connection or mechanism that enables two software applications or components to communicate with each other using a set of definitions and protocols. APIs are an accessible way to extract and share data within and across organizations.

APU: Refers to "accelerated processing unit," first created by AMD in 2011. An APU is a microprocessor that combines the central processing unit (CPU) with the graphics processing unit (GPU) on a single computer chip.

AR: Refers to "augmented reality," which is the integration of digital information with the user's environment in real time. Unlike VR, which creates a totally artificial environment, AR users experience a real world environment with generated perceptual information overlaid on top of it.

Arm: Arm Holdings plc, a legal entity. The acronym "Arm" stands for "Advanced RISC Machine." Arm uses an architecture based on RISC that is different from the once-standard x86 architecture built by Intel in the early days of computing that is based on CISC.

ASIC: Refers to "application-specific integrated circuit," which is a chip on which the pattern of connections has been set up exclusively for a specific function. Any chip that is custom-made is an ASIC, designed for a specific purpose. ASICs are typically designed and used by a single company in a specific system. They are expensive, time-consuming, and resource-intensive to develop, but they offer extremely high performance coupled with low power consumption.

ASSP: Refers to "application-specific standard processors." In contrast to ASICs, ASSPs are general-purpose devices intended for use by multiple system design houses.

CISC: Refers to "complex instruction set computing," which is distinct from RISC. RISC-V and Arm processors are based on RISC concepts and computing architectures while x86 processors from Intel and AMD employ CISC designs. CISC processors seek to minimize the number of instructions per program at the expense of the number of computer clock cycles per instruction. CISC computing seeks to complete a task in as few lines of code as possible, but a single CISC instruction may take multiple computer clock cycles to complete.

Both RISC and CISC can be seen as different "schools of thought about how to design a processor's ISA, or architecture for short.

Cloud Native Environment: Refers to the concept of building and running applications to take advantage of distributed computing offered by the cloud delivery model. Cloud native apps are designed and built to exploit the scale, elasticity, resiliency, and flexibility the cloud provides. In a cloud-enabled solution, a company's application is deployed in the public cloud but still requires a physical server of its own for operations. However, with a cloud-native system, the company does not require any computing infrastructure onsite.

CPU: Refers to "central processing unit," which is a chip that functions as the brains of computing devices. Also called the central processor, main processor or just processor, the CPU is made up of a set of electronic circuits that execute instructions comprising a computer program to perform basic arithmetic, logic, controlling, and I/O (input/output) operations specified in a computer program.

Compute: Compute means the processing of data.

Deep Learning: A subset of ML, deep learning is a computer science approach where neural networks (1) are trained to recognize patters from massive amounts of data (in the form of images, sounds and text), often better than humans, and in turn (2) provide predictions in a production process. Deep learning is often used to make non-linear, complex correlations, and needs to run on a specialized GPU.

Developers Toolkit: A developer's toolkit includes a range of tools and resources that enable third party semiconductor developers to create, analyze, test, and refine their designs, ensuring that the final products meet the required specifications and perform reliably in their intended applications. The toolkit might include design software, development kits and platforms, testing and measurement equipment, programming, manufacturing and prototyping tools, documentation and standards, and libraries and IP cores.

DPU: Refers to "data processing unit," which is a programmable computer processor designed to efficiently handle data-centric workloads, such as data transfer, reduction, security, compression, analytics and encryption, at scale in data centers. A DPU tightly integrates a general-purpose CPU with network interface hardware.

CPU vs. GPU chips

Edge: Refers to the "edge of a network." In a centralized network, client devices are connected to one server whose job is to process the information requested by the users and hand it back to them. While a centralized network is good enough to interact with simpler websites, companies have found that complex projects involving a great bank of customers work better with an edge server. Unlike a centralized network, because an edge server sits at the edge of a network, it improves latency, reduces loading times, and removes the load from the origin server. Instead of sending unprocessed data to the data center, edge servers process it themselves and send it back to the client machines.

EDA: Refers to "electronic design automation," a category of software tools for designing electronic systems such as ICs and printed circuit boards. According to EDA leader Synopsys, the goal of EDA is to assist in the definition, planning, design, implementation, verification, and subsequent manufacturing of semiconductor devices, or chips.

In the first area of focus of EDA (called technology computer-aided design, or TCAD), EDA tools are used to design and validate the semiconductor manufacturing process to ensure it delivers the required performance and density.

The second focus of EDA (called design for manufacturability, or DFM) EDA tools are used to verify that a design will meet all the requirements of the manufacturing process. Deficiencies in this area can cause the resultant chip to either not function or function at reduced capacity. There are also reliability risks.

In the third EDA focus area (called silicon lifecycle management, or SLM), after the chip is manufactured, there is a growing requirement to (1) monitor the performance of the device from post-manufacturing test to deployment in the field, (2) ensure the device continues to perform as expected throughout its lifetime and (3) ensure the device is not tampered with.

The fourth EDA focus area is called semiconductor intellectual property (or semiconductor IP). This market segment provides pre-designed circuits of varying complexity that may be used as-is or adapted for a particular application. Semiconductor IP allows highly complex chips to be designed in far less time since a lot of existing design work can be reused.

Embedded Processor: Refers to microprocessors that used in embedded systems such as electronic products ranging from video games to automotive control systems, industrial automation, sensors, signage and transportation, and home appliances, among other things. While these microprocessors are invisible and inaccessible to the end user, product designers use the computational capabilities of these embedded microprocessors to implement the operating features of electronic products and control systems. Both CISC and RISC microprocessors are used in the embedded market. Although the distinctions between RISC and CISC microprocessor technology have blurred considerably, CISC technology generally increases performance by using more complex instructions to achieve higher code density while RISC technology generally achieves substantial system performance and price/performance advantages by reducing the complexity of the processor instruction set, thereby enabling more effective use of pipelining and optimizing compilers, powerful hardware and software design techniques.

FPGA: A field-programmable gate array (FPGA) is a type of integrated circuit that can be programmed or reprogrammed after manufacturing. It consists of an array of programmable logic blocks and interconnects that can be configured to perform various digital functions. FPGAs are commonly used in applications where flexibility, speed, and parallel processing capabilities are required, such as in telecommunications, automotive, aerospace, and industrial sectors.

Global Foundries: Global Foundries, Inc., a Cayman Islands company, is the world's third largest semiconductor manufacturer and design company by revenue. Headquartered in in Malta, New York, Global Foundries was created by the the divestiture of the manufacturing arm of AMD. The company was privately owned by Mubadala Investment Company, the sovereign wealth fund of the United Arab Emirates, until until an an initial public offering in 2021. The company manufactures integrated circuits or wafers designed for markets such as smart mobile devices, automotive, aerospace and defense, consumer internet of things, and for data centers and communications infrastructure. It is the only fabrication company with operations in Singapore, the European Union, and the U.S.: one 200 mm and one 300 mm wafer fabrication plant in Singapore, one 300 mm plant in Dresden, Germany; one 200 mm plant in Burlington, Vermont. (where it is the largest private employer) and one 300 mm plant in Malta, New York.

GPU: Refers to "graphics processing unit," a type of microprocessor, capable of rendering graphics display on an electronic device, such as a gaming console or vehicle infotainment system. GPUs also have application for video editing and content creation, machine and deep learning, and artificial intelligence (AI). GPUs are also know as "video cards" or "graphics cards."

HPC: Refers to "high performance computing," which is the ability to process data and perform complex calculations at high speeds. To put it into perspective, a laptop or desktop with a 3 GHz processor can perform 3 billion calculations per second. While that is much faster than any human can achieve, it pales in comparison to HPC solutions that can perform quadrillions of calculations per second.

IC: An "integrated circuit" (also referred to as a chip or microchip) is a set of electronic circuits on one small flat piece of semiconductor material, usually silicon. Large numbers of miniaturized transistors and other electronic components are integrated together on the chip. This results in circuits that are orders of magnitude smaller, faster, and less expensive than those constructed of discrete components, allowing a large transistor count. The IC's mass production capability, reliability, and building-block approach to integrated circuit design have ensured the rapid adoption of standardized ICs in place of designs using discrete transistors.

These general-purpose devices that can be used for a variety of purposes and are found in everything from cell phones to cars to spacecraft. Unlike the microprocessor, which is designed to run software applications and gives better performance, integrated circuits are ideal for smaller applications that don't require as much computing power. Integrated circuits allow devices to be smaller and more lightweight. However, an integrated circuit cannot function independently because it doesn't have any program instructions or other information stored inside itself. In general, there are three types: (1) commodity integrated circuits, which are simple chips produced in large batches and used in single purpose appliances to perform repetitive processing routines (e.g., barcode scanners), (2) ICs made for a specific purpose, called application-specific integrated chip (ASIC) and (3) the all-in-one SoC, one of the newer types of chips.

Intel: Founded in 1968, Santa Clara, California-based Intel develops, manufactures, markets, and sells (1) platform products, such as central processing units and chipsets, and system-on-chip and multichip packages; and accelerators, boards and systems, connectivity products, and memory and storage products, (2) high-performance compute solutions for targeted verticals and embedded applications for retail, industrial, and healthcare markets; and solutions for assisted and autonomous driving comprising compute platforms, computer vision and machine learning-based sensing, mapping and localization, driving policy, and active sensors, and (3) workload-optimized platforms and related products for cloud service providers, enterprise and government, and communications service providers. Intel serves OEMs, original design manufacturer's, cloud service providers, and other equipment manufacturers.

IoT: Refers to "Internet of Things," which describes physical objects that are embedded with sensors, processing ability, and other technology that connects and exchanges data with other devices over the internet.

IP: Refers to the "intellectual property" of Arm, including (1) Arm architecture IP such as the company's RISC-based ISA, (2) physical IP, consisting of components used in the process of translating a design of complex SoC integrated circuits into a manufactured chip, (3) system IP consisting of components used in an SoC design (other than the GPU and CPU), such as interfaces to external components and on-dash chip interconnects that bridge functional components within the SoC, and (4) version extensions, updates, support, patches, and bug fixes.

ISA: Refers to "instruction set architecture," which specifies how software running on a CPU will behave. At the heart of every processor's functioning lies its ISA, which is often thought of as the blueprint that (1) outlines the set of instructions a processor can understand and execute and (2) shapes the capabilities and performance of a processor. The design of ISA influences how software is developed, and it has a lasting impact on a processor's efficiency, compatibility, and flexibility. ISAs can be broadly categorized into two types: open and closed. Closed ISAs, like Arm's, are proprietary and tightly controlled by their owner (e.g., Arm Holdings plc), offering established reliability and compatibility but limiting customization. On the other hand, open ISAs, exemplified by RISC-V, are community-driven and provide greater flexibility for customization, fostering innovation and adaptation to specific needs.

ISP: Refers to an "image signal processor" used to process images from sensors in digital cameras and other devices. It is often included in a SoC.

LLM: Refers to large language models. LLMs are advanced artificial intelligence systems designed to understand, generate, and interact with human language in a way that mimics human-like comprehension and responses. These models are trained on vast amounts of text data, enabling them to perform a wide range of language-related tasks, such as translation, summarization, answering questions, and creative writing. They work by analyzing the context and structure of language, learning patterns and nuances, and using this knowledge to generate coherent and contextually relevant text. Large language models, like OpenAI's GPT (Generative Pre-trained Transformer) series, represent the cutting-edge of natural language processing technology and are continually evolving to become more sophisticated and versatile in their language capabilities.

Machine Learning: Also known as "ML," machine learning is a subset of AI. ML is a computer science approach used in data science that leverages algorithms and data to learn how to make predictions or determinations. In other words, ML is an application of AI where machines/computer programs are given access to data and then use algorithms to find patterns in data (i.e., learn from the data). ML applications are often used to make simple, linear correlations, and can run on a CPU instead of a GPU that excels at running many smaller tasks at once.

With ML, a machine can automatically learn new information and improve itself from experiences without having to be programmed in a certain way since the machine will be able to teach itself new forms of information. However, the machines are not sentient (yet) and must be (1) told when they get something wrong or (2) given new data to figure out what the correct answer is. If new data or feedback is not given to the machine, it will not organically "learn" from its mistakes and adjust accordingly. As a result, ML does not provide "feedback"; it needs to be told explicitly how to fix a problem. AI, on the other hand, operates off a feedback loop.

Rather than having to input new data, AI is able to seek out new sources of data and rebuild itself based on those feedbacks, unlike ML.

Memory Chips: A memory chip is a semiconductor comprising numerous capacitors and transistors that can hold data (a) temporarily through random access memory (RAM) or (b) permanently through read only memory (ROM). Memory chips can be volatile (i.e., require power to maintain the stored information) or non-volatile (i.e., does not require a continuous power supply to retain the data or program code stored in a computing device).

Microcontrollers: A micro-controller is a small computer on a single VLSI integrated circuit. A microcontroller contains one or more CPUs (processor cores) along with memory, programmable input/output peripherals and program memory.

In contrast to the microprocessor (which is used in applications where (a) the task is not pre-defined but depends on the user or (b) intensive processing is required), a microcontroller is used for a specific task. Based on input given by the user or a sensor, the microcontroller performs a light amount of processing and gives the result as an output. In other words, microcontrollers are designed for embedded applications (such as washing machines and microwave ovens) where the task is pre-defined -- once the user sets the power and timing, the microwave provides cooked food; once the user sets the parameters, the washing machine provides clean clothes.

A microcontroller is similar to, but less sophisticated than, a SoC. An SoC may include a microcontroller as one of its components, but usually integrates it with advanced peripherals like a GPU, a Wi-Fi module, or one or more co-processors.

Microprocessors: Micro-processors are used in (and mounted on the mother-boards of) virtually all desktop and laptop computers and other types of electronic equipment, such as cell phones, tablets, and components of cars. Microprocessors require extremely precise manufacturing to function properly, are power-hungry, and are better suited for more demanding tasks. They respond to user input, process instructions and perform the calculations and logic operations used to run software programs—i.e., play games, browse websites, edit photos, create documents, stream media, and perform mathematical calculations. The microprocessor contains all the components found in an IC — a CPU, input / output ports, and memory — but it also includes its own non-volatile storage for programs and data. This allows users to run software on computers without needing additional hardware devices.

Middleware: Refers to software that different applications use to communicate with each other. Middleware acts as a bridge between diverse technologies, tools, and databases so the user can integrate them seamlessly into a single system.

Mobile Processors: A mobile processor is a microprocessor designed for mobile devices such as laptops and cellphones. They are the brains of the device. Mobile processors tend to be less powerful than PC processors due to their lower TDP and power consumption. PC processors are optimized for high-performance, while mobile processors are optimized for low power consumption and long battery life. Mobile processors are typically cheaper than PC processors which is one of the reasons mobile devise are more affordable than PCs. It is worth noting that there are some high-end mobile processors that are on par with some PC processors in terms of performance, but these are usually the exception and not the rule.

NPU: Refers to "neural processing unit," also known as a neural processor, tensor processing unit, or AI accelerator). The NPU is a specialized processor that implements all necessary control and arithmetic logic necessary to execute ML algorithms, typically by operating on predictive models such as artificial neural networks or random forests. In short, NPUs, such as the M2 chip used in latest versions of Apple's laptops and tablets, specialize in the acceleration of ML algorithms (10,000 times less time than a GPU).

NVIDIA: NVIDIA Corporation is the worldwide leader in visual computing technologies and the inventor of the graphic processing unit (GPU). Over the years, the company's focus has evolved from PC graphics to artificial intelligence (AI) based solutions that now support high performance computing (HPC), gaming and virtual reality (VR) platforms. NVIDIA's GPU success can be attributed to its parallel processing capabilities supported by thousands of computing cores, which are necessary to run deep learning algorithms. The company's GPU platforms are playing a major role in developing multi-billion-dollar end-markets like robotics and self-driving vehicles. NVIDIA is also a dominant name in the data center, professional visualization and gaming markets where Intel and Advanced Micro Devices are playing catch-up. NVIDIA's partnership with almost all major cloud service providers and server vendors is the catalyst.

PCIe: Refers to "peripheral component interconnect express," an interface standard for connecting high-speed components. Every desktop PC motherboard has a number of PCIe slots one can use to add GPUs, RAID (redundant array of independent disks) cards, Wi-Fi cards or SSD (solid-state drive) add-on cards.

RF: Refers to "radio frequency," which is used in various electronics and appliances (e.g., radios and microwave ovens), TV broadcasting, cellular telephones, satellite communications, radars, and industrial heaters and sealers. RF is the oscillation rate of an electromagnetic field.

RISC: Refers to "reduced instruction set computing" architectures and principles on which Arm's proprietary ISA is based. A RISC architecture has simple instructions that can be executed extremely quickly in a single computer clock cycle. Arm's RISC-based ISA has become the dominant choice for embedded systems and mobile devices due to its long-standing presence in the market, as well as years of trust and expertise cultivated, resulting in a widespread positive reputation.

Semiconductors: Also known as chips or processors, semiconductors are made by imprinting a network of electronic circuits/components onto a wafer that partially conducts electricity, which can then perform various functions—e.g., (a) processing, amplifying and selectively filtering electronic signals (b) controlling, electronic system functions, and (c) storing and transmitting data.

Semiconductor Foundries: Also known as fabs, foundries are the facilities where semiconductor chips or devices are manufactured. These highly complex and expensive facilities are either owned and operated by large, vertically integrated semiconductor companies (e.g., Infineon, Samsung, Intel, and Renesas) or operated as independent, "pure-play" manufacturing service providers (e.g., TSMC or Global Foundries)—i.e., the dominant business model.

Smart NIC: Refers to "smart network interface card," which is a programable accelerator that makes data center networking, security and storage efficient and flexible. See DPU above.

SoC: Refers to "system on a chip" or system-on-chip, which are among the newest types of chips. SoCs integrate or build into one chip most or all electronic components needed for an entire electronic system, such as an on-chip central processing unit (CPU), memory interfaces (RAM or random-access memory), input/output ports, input/output interfaces, and secondary storage interfaces, often alongside other components such as radio modems and GPU — all on a single substrate or microchip. SoCs may contain digital, and also analog, mixed-signal, and often radio frequency signal processing functions. SoCs break from the traditional approach to system architecture—e.g., motherboards on which each component is discretely installed. This enables the creation of smaller and more efficient devices. For this reason, SoCs are driving innovation in the creation of slim net-books, lap-tops, smart phones, and IoT devices.

Transistors: A miniature semiconductor that regulates or controls current or voltage flow in addition to amplifying and generating these electrical signals and acting as a switch/gage for them.

VR: Refers to "virtual Reality"; VR is distinct from AR because VR creates a totally artificial world.

VLSI: Means "very-large-scale-integration"--creating an integrated circuit by combining thousands of transistors into a single chip.

WLAN: Refers to "wireless local area network," which is a wireless distribution method for two or more devices; it uses high frequency radio waves and includes an access point to the Internet.

x86: x86 is a family of complex ISAs (or CISCs) initially developed by Intel. x86 processors from AMD and Intel dominate in computers and servers, while Arm RISC-based processors from Apple and Qualcomm dominate in tablets and smartphones. Arm-based processors have also been strong in embedded applications. That strength is being challenged by RISC-V, which is increasingly being used in embedded systems, automotive systems, disc drives, AI, and ML applications.

3G, 4G, and 5G: 3rd, 4th , and 5th generations of wireless, mobile, broadband / cellular network technology. The latest generation is designed to deliver higher multi-Gbps peak data speed, ultra-low latency, more reliability, massive network capacity, increased availability, and a more uniform user experience than all previous generations.

OVERVIEW OF SEMICONDUCTORS INDUSTRY: ESSENTIAL FOR MODERN LIVING

September 2025

WHAT SEMICONDUCTORS DO

Overview: Semiconductors are an essential component of modern life! They are tiny electronic circuits made up of materials that partially conduct electricity. A broad portfolio of semiconductors are essential components of electronic devices and enable these devices to perform their functions. As such, semiconductors are considered the brains for millions of electronic devices and electrical systems located inside space vehicles, TVs, radios, video game consoles, car computers and electrical systems, smartphones, personal computers, supercomputers, advanced medical diagnostic equipment, appliances, military weapons systems, renewable energy systems, industrial equipment, data center servers, and telecommunications networks.

Composition and Function: A semiconductor is a material that allows current to flow more than an insulator but less than a pure conductor. Memory chips, microprocessors, and integrated chips are some examples of semiconductors. Semiconductors are made by imprinting a network of electronic components onto a semiconductor wafer. These devices are designed to perform various functions such as processing, amplifying and selectively filtering signals, controlling electronic system functions and processing, and transmitting and storing data.

Products and Devices That Incorporate Semiconductors

GENERAL APPLICATIONS

- Vehicles (See Specific Cases)
- Space Vehicles
- Railcars
- Robotic Welding Machine
- Video Game Consoles
- Medical Diagnostic Equipment
- Data Center Servers
- Embedded Systems (Small Computers that Form Part of a Larger Machine; They Control the Device and Allow User Interaction)
  - GPS Systems
  - TVs
  - Smart Watches
  - Fitness Trackers
  - Engine Management Systems
  - Central Heating Systems
- IT Devices (see Specific Cases)
- Factory Automation Systems
- Appliances
- Smart Phones
- Supercomputers
- Oil, Gas and NGL Pipelines
- Military Weapons Systems
- Thermoelectric Applications (Including Lighting and LED Displays)
- Radios
- Telecommunications Networks
- Networking Equipment
- Solar Cells and Other Renewable Energy Systems

SPECIFIC CASES

Vehicles

The typical modern car is loaded with over 1,400 semiconductors to control safety, powertrain, electrical, comfort, infotainment, and connectivity systems. For this reason, vehicles are often referred to as "computers on wheels."

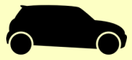

Connectivity
- Controller Area Network (CAN)
- Broadband, Wi-Fi, Bluetooth
- Over-the-Air Software Updates

Comfort
- Window/Mirror Controls
- Vibration Control
- Seat Controls
- Climate Controls
- Noise Suppression
- Adaptive Cruise Control

Infotainment
- Audio/Video
- Driver Display
- Navigation

Powertrain
- Engine Control
- Fuel-Injection System
- Hybrid-Electric Control
- Transmission Control
- Electric Power Steering

Safety
- Airbag Controls
- Collision Avoidance
- Parking-Assist
- Power Locks
- Braking-Assist
- Tire-Pressure-Monitoring
- Traction-Control System

Electrical
- Starter
- Lighter
- Vehicle-Diagnostics

IT Devices

Most people associate semiconductors with information technology devices, such as personal computers (desktop and laptop computers), mobile devices (smart phones and tablets), and servers in data centers.

Personal Computers

In 2021, 342 million PCs were shipped around the world. It was expected that 310 million PCs were to be shipped in 2022. However, PC shipments only reached 286.2 million units in 2022, a 16.2% decrease from the previous year. Most PCs include only one chip.

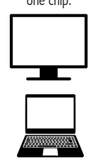

Mobile Devices

There are 6.7 billion smartphone subscriptions globally (85% of the world population). By 2028, the rise in the number of smartphone subscriptions is projected to grow to 7.7 billion. Current generations of smartphones come in a range of multi-range of processors, such as Dual-Core (2), Quad-Core (4), and Octa-Core (8), with the Octa-Core processor (used by Apple in iPhone 13 and 14) being the most powerful.

Servers

Servers today feature one or two x86 chips, or maybe an Arm processor. In 5 or 10 years they will feature many more. The typical data center with 850 Megawatts of capacity can run thousands to millions of servers:

- 6,314,256 Low-Powered 1U Servers
- 1,768,000 Mid-Powered 1U Servers
- 803,608 High-Powered 1U Servers

OVERVIEW OF SEMICONDUCTORS: CHIPMAKERS, PRODUCTION PROCESS, TYPES, AND APPLICATIONS

September 2025

SEIMCONDUCTOR ECOSYSTEM

INDUSTRY PARTICIPANTS

Key semiconductor players include:

1. **DESIGN ARCHITECTS** (Design architects research and develop instruction set architectures for processors as well as physical and system IP, which they license to system companies for customization and manufacturing of chips).
   arm
2. **SYSTEM COMPANIES** (Semiconductor companies whose research and integrated circuit designs set the stage for chip production).
   AMD, Qualcomm, BROADCOM
3. **FABS OR FOUNDRIES** (Pure-play foundries and integrated circuit manufacturers perform front-end wafer manufacturing operations. Foundry operators may be System Companies or a third-party contract manufacturers (if the System Company operates on a "fabless" basis)).
   tsmc, Infineon, intel
4. **ELECTRONICS MANUFACTURING SERVICES** (EMS) (Companies provide integrated circuit assembly, packaging, testing, and repair services to the System Company. These companies are also capable of designing and manufacturing semiconductors).
   Amkor Technology, FOXCONN, JABIL, ASE GROUP
5. **OEM CONSUMERS** (OEM consumers are the original equipment manufacturers of electronic devices (or the makers of electronic components and assemblies of such devices) into which finished chips will be incorporated).
   DELL Technologies, verizon, aws, Ford, RTX

PRODUCTION PROCESS

According to the Semiconductor Industry Association, the semiconductor production process consists of nine steps—research, design, raw materials procurement, ingot production, blank wafer production, finished wafer production, finished wafer cutting, finished wafer packaging, and embedding finished wafers in circuit boards.

FABLESS SYSTEM COMPANY → **FOUNDRY** (May partner with System Company on Design and R&D) → **EMS COMPANY** → **OEM CONSUMER**

FABLESS MODEL

1. **RESEARCH** — Pre-competitive research is conducted.
2. **DESIGN** — Engineers use highly sophisticated equipment to design semiconductors, similar to how architects design buildings.
3. **RAW MATERIALS** — Many semiconductors start out as germanium or sand, which contains a large amount of silicon, or compounds such as gallium arsenide. In a process called doping, small amounts of impurities are added to these pure elements or compounds, causing large changes in the conductivity of the material.
4. **INGOT** — The sand is purified and melted into solid cylinders called ingots, weighing up to 200+ lbs.
5. **BLANK WAFER** — The ingot is then sliced into very thin (1mm) silicon discs and polished to a flawless finish.
6. **FINISHED WAFER** — Next, wafers are printed with highly intricate circuit designs that will later become individual chips.
7. **CUT WAFER** — The silicon wafer containing finished semiconductors, sometimes as many as 70,000 per wafer, is then cut up into tiny individual semiconductors called dies.
8. **PACKAGED CHIP** — The dies are then packaged into finished semiconductors, which can then be placed into electronic devices.

Fab-less Model

8A. **DELIVERY CHIPS TO FABLESS SYSTEM COMPANY** — The packaged chips are delivered to the Fabless System Company.

Fabless Model

9. **CHIP ON CIRCUIT BOARD** — Finished semiconductors are embedded in countless electronics devices, from personal computers, servers and smartphones to highly advanced medical equipment, vehicle electrical systems, and supercomputers.

Foundry Model

FOUNDRY MODEL

FOUNDRY → **FOUNDRY END CONSUMER**

CHIPMAKRERS

SYSTEM COMPANIES

The numerous chipmakers range from (1) household names with global reach to (2) smaller suppliers little known outside their specialized market niche.

Semiconductor firms focus their operations around one of three models:

6. **FABLESS** (Fabless companies focus only on chip design and then outsource the manufacturing of their chip designs to third party producers).
   Qualcomm, NVIDIA.
7. **INTEGRATED DVICE MANUFACTURERS** (IDMs) (IDMs conduct both design and manufacturing activities).
   SAMSUNG, intel.
8. **FOUNDRIES** (Pure-play foundries are companies that focus primarily on manufacturing for third parties.
   GlobalFoundries, tsmc

DESIGN ARCHITECTS

arm, RISC-V, intel x86

FABLESS COMPANIES

Qualcomm, HUAWEI, SONY, TOSHIBA, HISILICON (Huawei-Owned), NVIDIA, AMD, NXP, BROADCOM, MEDIATEK, Apple-Designed, TSMC-Produced

INTEGRATED DEVICE MANUFACTURERS

FUJITSU, SAMSUNG, TEXAS INSTRUMENTS, intel, Infineon, SK hynix, MICRON, RENESAS

PURE-PLAY FOUNDRIES

GlobalFoundries, tsmc, UMC

CHIP TYPES

KEY CHIP CATEGORIES

Semiconductors can be divided into three main categories:

9. MEMORY CHIPS
10. PROCESSORS AND CONTROLLERS
11. INTEGRATED CIRCUITS (Including Accelerated Processing Units (APUs), Neural Processing Units (NPUs), Data Processing Units (DPUs), Graphics Processing units (GPUs), and Systems on Chip (SOCs))

MEMORY CHIPS

MEMORY CHIPS

A memory chip is a semiconductor comprising numerous capacitors and transistors that can hold data (a) temporarily through random access memory (RAM) or (b) permanently through read only memory (ROM). Memory chips can be volatile (i.e., require power to maintain the stored information) or non-volatile (i.e., does not require a continuous power supply to retain the data or program code stored in a computing device).

PROCESSORS AND CONTROLLERS

CENTRAL PROCESSING UNITS (CPUs)

The CPU is a chip that functions as the brains of the personal computer (PC). Also called the central processor, main processor or just process, the CPU is made up of a set of electronic circuits that execute instructions comprising a computer program to perform basic arithmetic, logic, controlling and I/O (input/output) operations specified in computer program.

MICROPROCESSORS

Microprocessors are used in (and mounted on the mother-boards of) virtually all desktop and laptop computers. They are also found in other types of electronic equipment, such as cell phones, tablets, and cars. Microprocessors require extremely precise manufacturing to function properly, are power-hungry, and are better suited for more demanding tasks. They respond to user input, process instructions and perform the calculations and logic operations used to run software programs—i.e., play games, browse websites, edit photos, create documents, stream media, and perform mathematical calculations. The microprocessor contains all the components found in an integrated circuit — a CPU, input / output ports, and memory — but it also includes its own non-volatile storage for programs and data. This allows users to run software on computers without needing additional hardware devices.

MICROCONTROLLERS

In contrast to the microprocessor (which are used in applications where (a) the task is not pre-defined but depends on the user or (b) intensive processing is required), a microcontroller is used for a specific task. Based on input given by the user or a sensor, the microcontroller performs a light amount of processing and gives the result as an output. Examples include the washing machine and microwave oven, where the task is pre-defined -- once the user sets the power and timing, the microwave gives the user cooked food; once the user sets the parameters, the washing machine gives the user clean clothes.

INTEGRATED CIRCUITS

INTEGRATED CIRCUITS (ICs)

Integrated circuits are general-purpose devices that can be used for a variety of purposes and are found in everything from cell phones to cars to spacecraft. Unlike the microprocessor, which is designed to run software applications and gives better performance, integrated circuits are ideal for smaller applications that don't require as much computing power. Integrated circuits allow devices to be smaller and more lightweight. However, an integrated circuit cannot function independently because it doesn't have any program instructions or other information stored inside itself. In general, there are three types: (1) commodity integrated circuits (CICs), which are simple chips produced in large batches and used in single purpose appliances to perform repetitive processing routines (e.g., barcode scanners), (2) ICs made for a specific purpose, called application-specific integrated chip (ASIC) and (3) the all-in-one "system on a chip" (SoC), one of the newer types of chips.

ACCELERATED PROCESSING UNITS (APUs)

Created by AMD in 2011, an accelerated processing unit (APU) is a microprocessor that combines the central processing unit (CPU) with the graphics processing unit (GPU) on a single computer chip.

GRAPHICS PROCESSING UNITS (GPUs)

A type of microprocessor, a graphics processing unit (GPU) is capable of rendering graphics display on an electronic device, such as a gaming console, personal computer or vehicle infotainment system. GPUs also have application for video editing and content creation, machine and deep learning, and artificial intelligence (AI).

NEURAL PROCESSING UNITS (NPUs)

A neural processing unit (NPU), such as the M2 chip used in latest versions of Apple's laptops and tablets, is a microprocessor that specializes in the acceleration of machine learning algorithms (10,000 times less time than a GPU), typically by operating on predictive models such as artificial neural networks (ANNs) or random forests (RFs). It is also known as neural processor.

DATA PROCESSING UNITS (DPUs)

DPUs are the new kids on the block. A DPU is a programmable processor designed to efficiently handle data-centric workloads such as data transfer, reduction, security, compression, analytics, and encryption, at scale in data centers. DPUs work with CPUs and GPUs to enhance the computing power and handling of increasingly complex modern data workloads. The DPU market has gained traction due to the rise in demand for AI, machine learning, deep learning, Internet of Things, 5G, and complex cloud architectures.

SYSTEMS ON A CHIP (SoCs)

A form of integrated circuit, the SoC is among the newest types of chip. All the electronic components needed for an entire system (e.g., CPU, Random Access Memory (or RAM), storage, I/O (input/out ports) are built into a single chip. SoCs break from the traditional approach to system architecture (e.g., with motherboards) where each component is discretely installed. This enables the creation of smaller and more efficient devices. For this reason, SoCs are driving innovation in the creation of netbooks, laptops, smartphones, Internet of Things (IoT) devices.