
AMD
together we advance_data centers

Driving Significant Growth

Setting the Foundation 2015	"Zen" Launch 2017	High-Performance Leadership 2021	High-Performance and Adaptive Leader 2023
Intel: $150B	Intel: $171B	Nvidia: $444B	Nvidia: $1,103B
Qualcomm: $105B	Broadcom: $102B	Intel: $231B	Broadcom: $369B
Broadcom: $37B	Nvidia: $89B	Broadcom: $191B	**AMD $184B (#3)**
Mediatek: $21B	Qualcomm: $85B	Qualcomm: $152B	Intel: $149B
Nvidia: $11B	Mediatek - $13B	**AMD $97B (#7)**	Qualcomm: $128B
Marvell: $7B	**AMD $11B (#15)**	Mediatek: $56B	Marvell: $52B
AMD $2B (#18)	Marvell: $9B	Marvell: $44B	Mediatek: $34B

AMD is Now World's Third Most Valuable Semiconductor Company

Growing Market Momentum

AMD EPYC™ CPUs Trusted to Power ~40% of the World's Servers

>18x Server CPU Market Share Growth

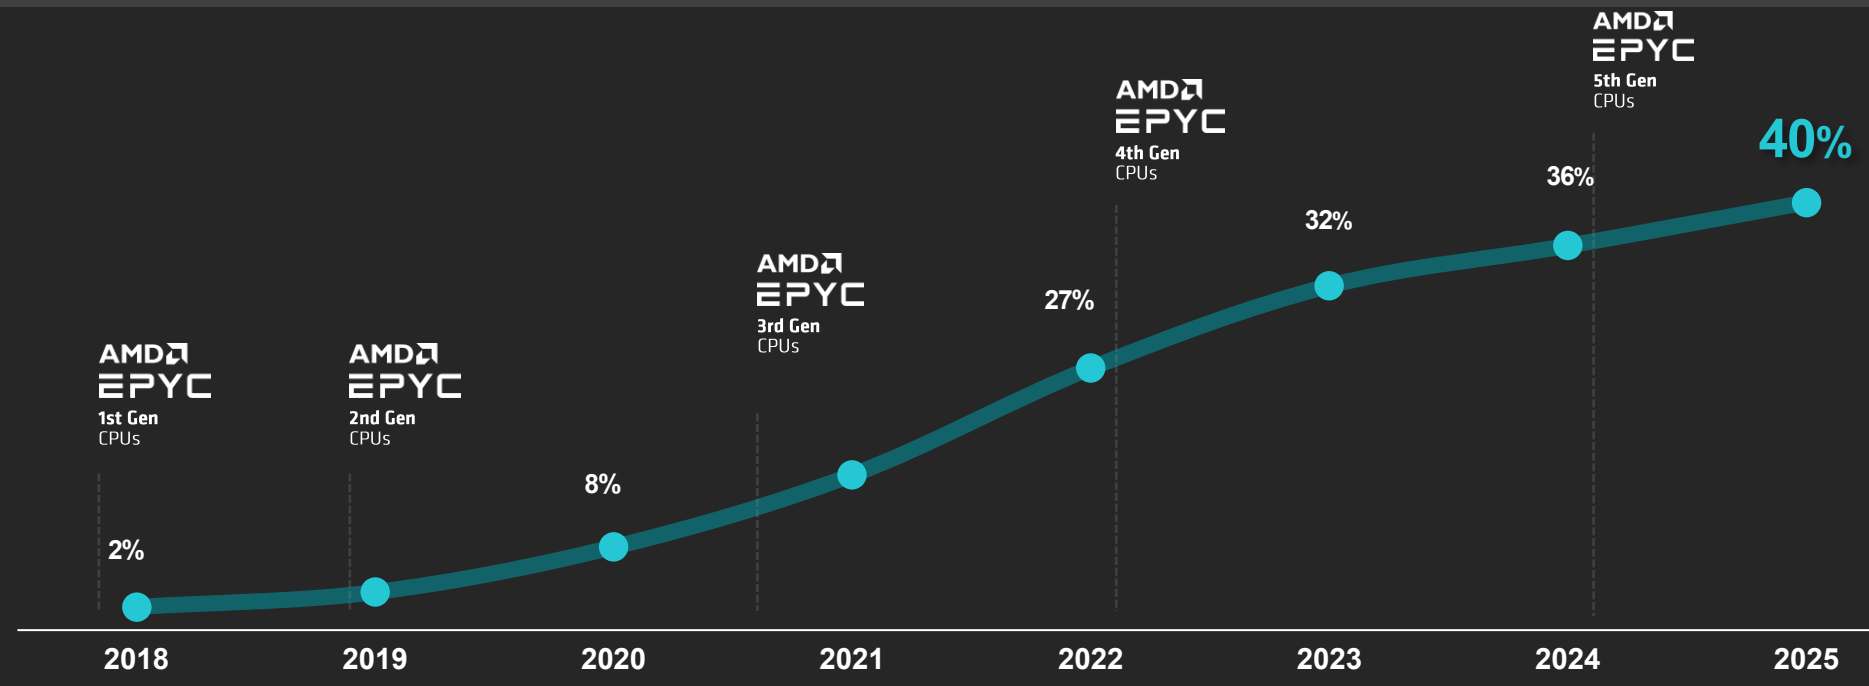

Digital | NETFLIX | Uber | Meta | zoom

Enterprise | BEST BUY | IBM | Emirates NBD | NISSAN

Perfectly Positioned to Win More Together

Source: Mercury Research Sell-in Revenue Shipment Estimates, 2025 Q1

ADDRESSING THE INDUSTRY'S BIGGEST CHALLENGES

MODERNIZING THE DATACENTER WITH AI

- Finding the Budget to Invest
- Physical Space in Datacenter
- Not Enough Power Available
- AI Increasing Power Challenge

DISRUPTION TO VMWARE LICENSING COSTS

- Broadcom Acquisition of VMWare
- New Licensing Model Announced
- Shift to Core-Based Licensing
- Significant Increase in Expected Licensing Costs to Enterprise
- ~2-10X increase in Licensing costs

CORPORATE SUSTAINABILITY

- Half of World's Largest Companies Committing to Net Zero
- Rising Electricity Costs -> California electricity up 45.9% since 2020

AMD VALUE PROPOSITION

PERFORMANCE

The Highest Processor Performance and Scalability

SECURITY

Advanced Security Features Help Mitigate Risk

COST OF OWNERSHIP

Capture the Full Benefits of Your IT Investments

WITH UNMATCHED SYSTEM ARCHITECTURE TO POWER THE EXASCALE ERA FROM SUPERCOMPUTERS TO CLOUD

Fifth Generation of CPU Leadership

6x Core Count
Across 5 Generations

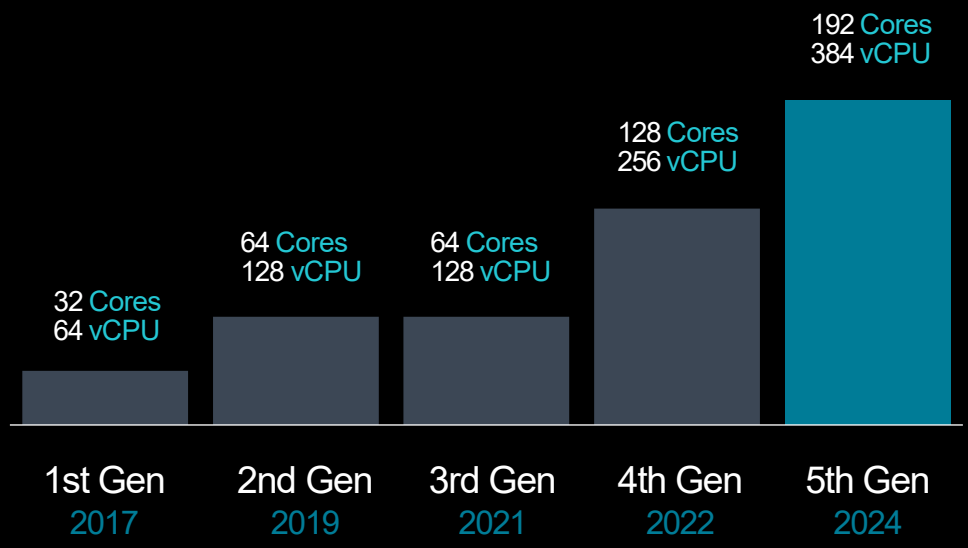

10.9x Performance
Across 5 Generations

Relative SPECrate2017_int_base best score by end of launch year

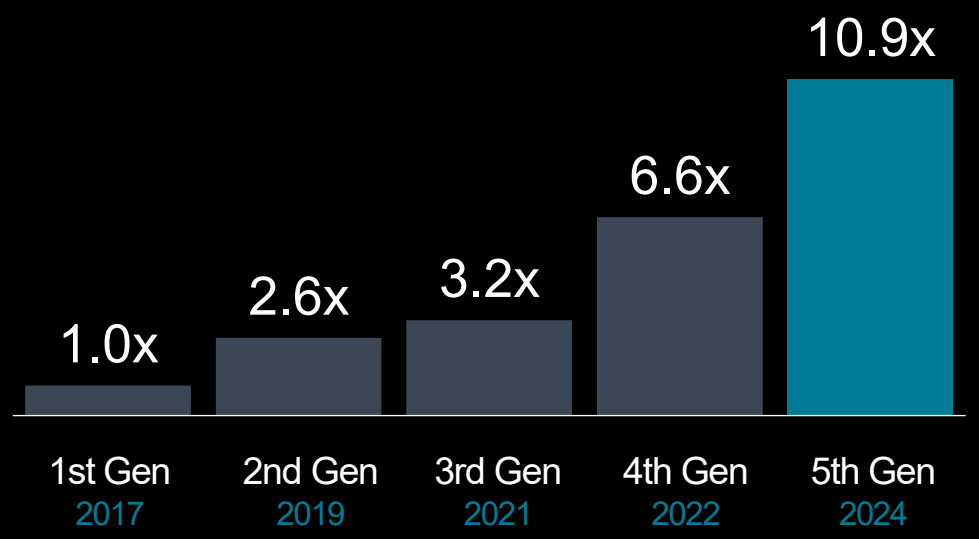

As of 10/10/2024. See endnote 9xx5-069A

AMD EPYC

THE BEST GETS EVEN BETTER

300+ WORLD RECORDS AND *STILL* COUNTING

BUSINESS APPS

8 ERM/SCM Business

48 Enterprise server-side Java®

18 Enterprise Energy Efficiency

4 FinTech

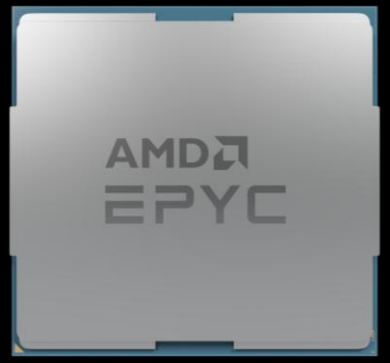

DATA MANAGEMENT & ANALYTICS

16 Business Intelligence

6 Structured Database Management

2 Unstructured Analytics Database

9 AI Platform/Text & Media Analytics

INFRASTRUCTURE/HCI/SDI

23 Cloud and Virtualization

15 Integer Performance/General Purpose

8 Integer/General-Purpose Energy Efficiency

3 VM Energy Efficiency

ENGINEERING/TECHNICAL

73 Massively Parallel applications

37 Modeling & Simulation

16 Floating Point Compute Intensive apps

12 HPC Energy Efficiency

CONTENT APPS

5 Rendering

DIGITAL SERVICES

1 Social Networking BI

AMD EPYC Family of Processors as of 11/10/22, see amd.com/worldrecords for the full list

Why it matters

Of recent survey respondents*

- *70% of technology sourcing, procurement and vendor management leaders will have environmental-sustainability-aligned performance objectives for their functions, by 2026.*
- *75% of organizations will increase business with IT vendors that have demonstrable sustainability goals and timelines and will seek to replace vendors that do not, by 2026.*

AMD embraces the responsibility to protect our planet and the opportunity to help others reduce IT-related energy use and greenhouse gas (GHG) emissions.

* Gartner®: Predicts 2023: Environmental Sustainability is Now an IT Sourcing Imperative – January 2023

Advancing data center sustainability

The AMD “30x25” goal is to deliver 30x more energy efficiency for our accelerated compute nodes powering servers for AI-training and HPC (2020-2025).[2] The goal represents:

- 2.5x acceleration of the industry trends from 2015-2020 (measured by worldwide energy consumption for these computing segments)
- 97% reduction in energy use per computation from 2020-2025

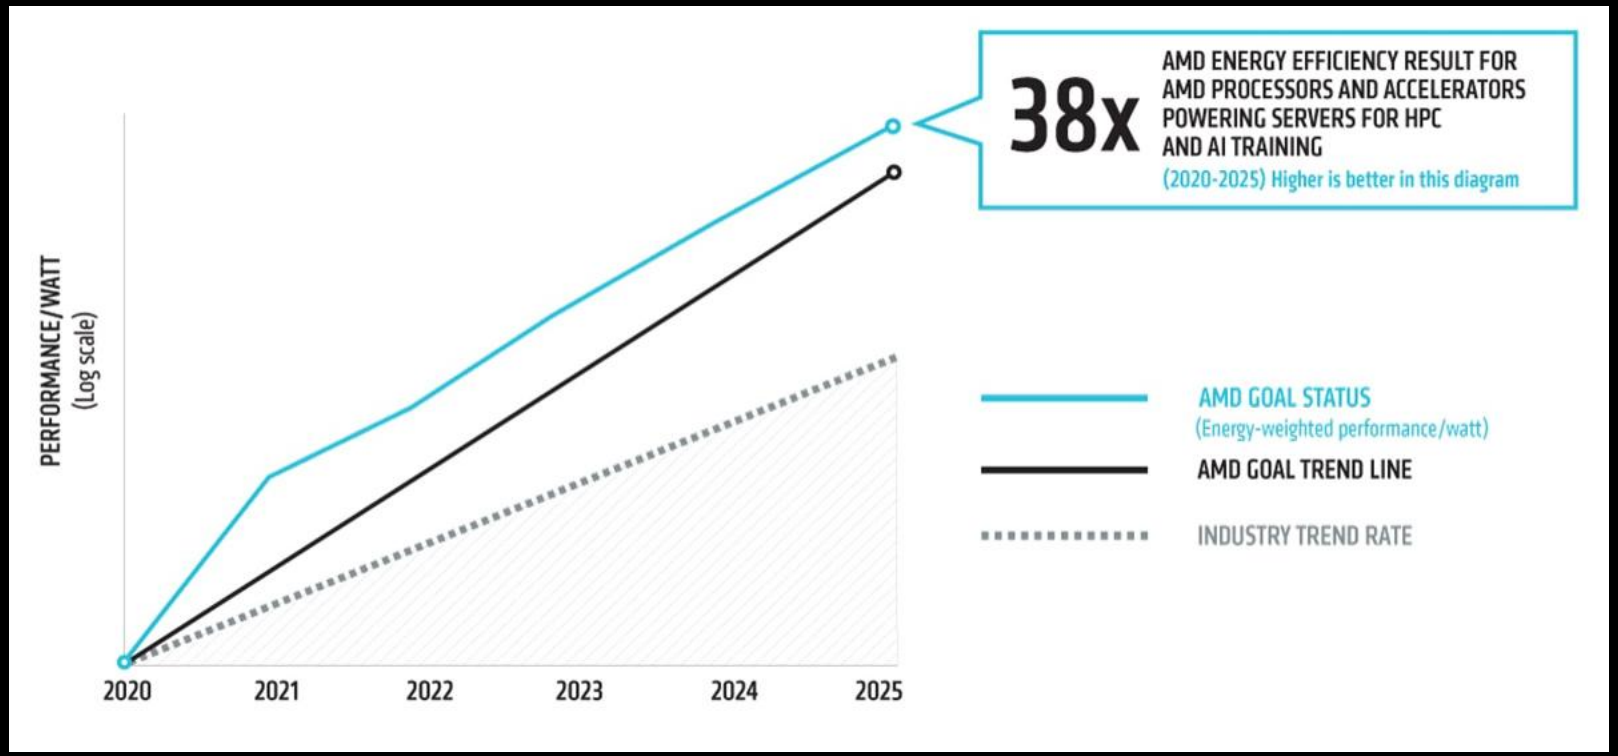

Learn more at https://www.amd.com/en/corporate-responsibility/data-center-sustainability

In the past 5 years, we've increased our own data center performance over 8x while using only 3x the number of compute CPU cores.

And we achieved that while maintaining the same datacenter footprint over the last two years despite the huge increase in compute capability.

Pushing the envelope in performance and energy efficiency is a foundational part of our product development strategy.

8x

increase in **AMD's datacenter performance** since 2018

3x

increase in **AMD's datacenter number of cores** since 2018

The Exascale Era

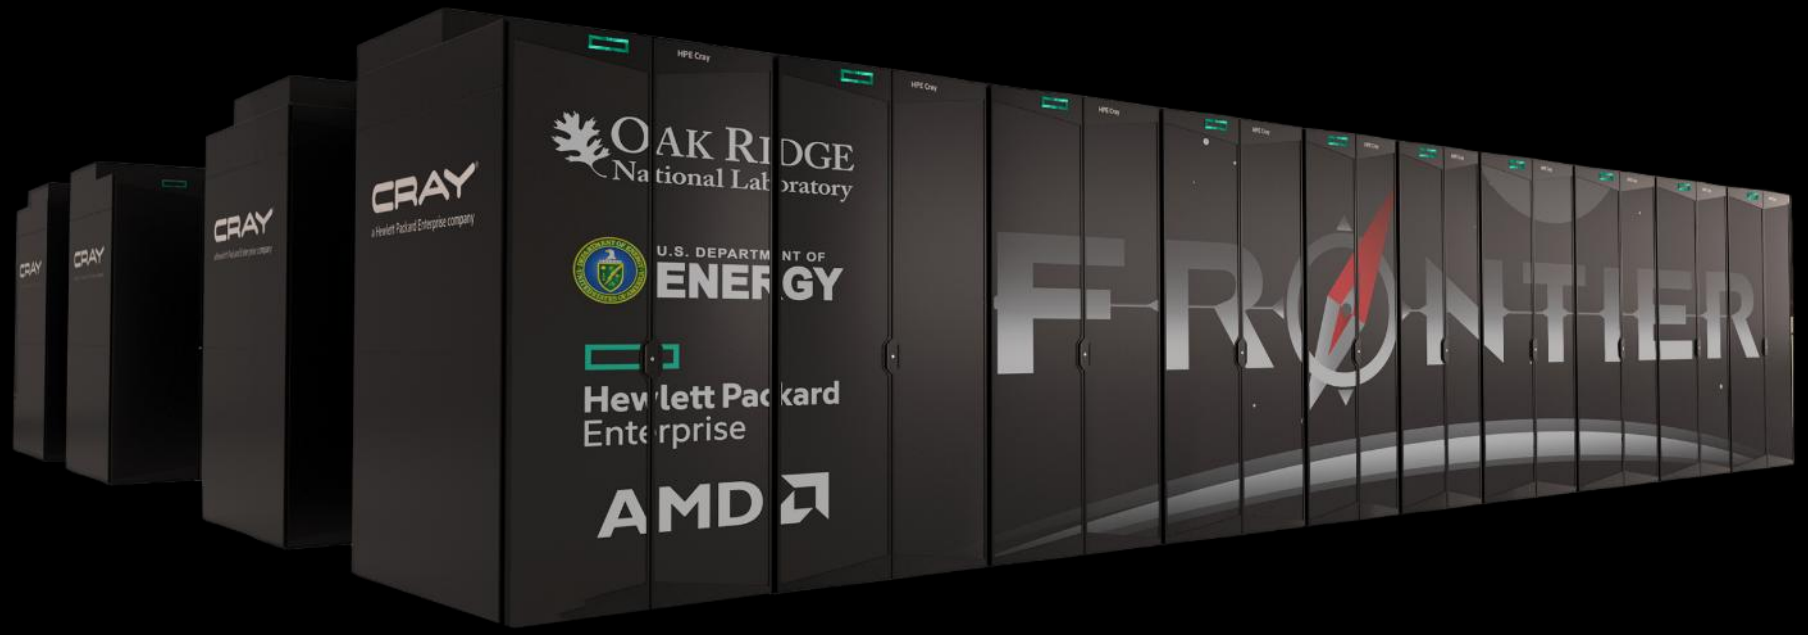

World's 1st Exascale Supercomputer also World's Most Energy Efficient.

Faster than the next best 7 supercomputers combined

9,408 CPUs + 37,632 GPUs connected by 90 miles of network cables

Sources: https://www.top500.org/system/180047/
https://top500.org/news/ornls-frontier-first-to-break-the-exaflop-ceiling/

AMD EPYC™ CPUs

Workload optimized platforms

NO COMPROMISE

ONE SOCKET

Right-Size Systems to Workloads
without Compromise

Outstanding Performance | Highly Scalable | Cost Optimized

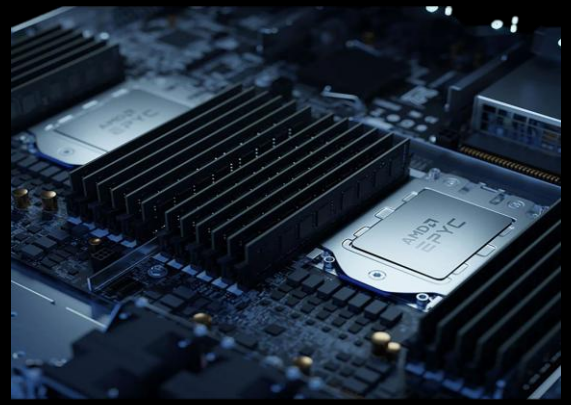

LEADERSHIP

TWO SOCKET

High-Performance, Balanced Architecture

Exceptional Performance | Maximum Scalability | Cost Effective

AMD EPYC™ 9004 Series Processor

- All-in Feature Set support
- 12 Channels of DDR5-4800
- Up to 6TB DDR5 memory capacity
- 128 lanes PCIe® 5
- 64 lanes CXL 1.1+
- AVX-512 ISA, SMT & core frequency boost
- AMD Infinity Fabric™
- AMD Infinity Guard

Cores	AMD EPYC	Base/Boost* (up to GHz)	Default TDP (w)	cTDP (w)
96 cores	9654/P	2.40/3.70	360w	320-400w
84 cores	9634	2.25/3.70	290w	240-300w
64 cores	9554/P	3.10/3.75	360w	320-400w
64 cores	9534	2.45/3.70	280w	240-300w
48 cores	→ 9474F	3.60/4.10	360w	320-400w
	9454/P	2.75/3.80	290w	240-300w
32 cores	→ 9374F	3.85/4.30	320w	320-400w
32 cores	9354/P	3.25/3.80	280w	240-300w
32 cores	9334	2.70/3.90	210w	200-240w
24 cores	→ 9274F	4.05/4.30	320w	320-400w
	9254	2.90/4.15	200w	200-240w
	9224	2.50/3.70	200w	200-240w
16 cores	→ 9174F	4.10/4.40	320w	320-400w
	9124	3.00/3.70	200w	200-240w

*SEE ENDNOTES: EPYC-018, GD-183

AMD EPYC™ 9005 Series Processors

- Increased core density
- Energy efficient
- Broad OPN stack

Cores	AMD EPYC	CCD (Zen5/Zen5c)	Base/Boost* (up to GHz)	Default TDP (W)	L3 Cache (MB)	Price (1 KU, USD)
192 cores	9965	"Zen5c"	2.25 / 3.7	500W	384	$14,813
160 cores	9845	"Zen5c"	2.1 / 3.7	390W	320	$13,564
144 cores	9825	"Zen5c"	2.2 / 3.7	390W	384	$13,006
128 cores	9755 9745	"Zen5" "Zen5c"	2.7 / 4.1 2.4 / 3.7	500W 400W	512 256	$12,984 $12,141
96 cores	9655 9655P 9645	"Zen5" "Zen5" "Zen5c"	2.6 / 4.5 2.6 / 4.5 2.3 / 3.7	400W 400W 320W	384 384 256	$11,852 $10,811 $11,048
72 cores	9565	"Zen5"	3.15 / 4.3	400W	384	$10,486
64 cores	9575F 9555 9555P 9535	"Zen5" "Zen5" "Zen5" "Zen5"	3.3 / 5.0 3.2 / 4.4 3.2 / 4.4 2.4 / 4.3	400W 360W 360W 300W	256 256 256 256	$11,791 $9,826 $7,983 $8,992
48 cores	9475F 9455 9455P	"Zen5" "Zen5" "Zen5"	3.65 / 4.8 3.15 / 4.4 3.15 / 4.4	400W 300W 300W	256 256 256	$7,592 $5,412 $4,819
36 cores	9365	"Zen5"	3.4 / 4.3	300W	192	$4,341
32 cores	9375F 9355 9355P 9335	"Zen5" "Zen5" "Zen5" "Zen5"	3.8 / 4.8 3.55 / 4.4 3.55 / 4.4 3.0 / 4.4	320W 280W 280W 210W	256 256 256 128	$5,306 $3,694 $2,998 $3,178
24 cores	9275F 9255	"Zen5" "Zen5"	4.1 / 4.8 3.25 / 4.3	320W 200W	256 128	$3,439 $2,495
16 cores	9175F 9135 9115	"Zen5" "Zen5" "Zen5"	4.2 / 5.0 3.65 / 4.3 2.6 / 4.1	320W 200W 125W	512 64 64	$4,256 $1,214 $726
8 cores	9015	"Zen5"	3.6 / 4.1	125W	64	$527

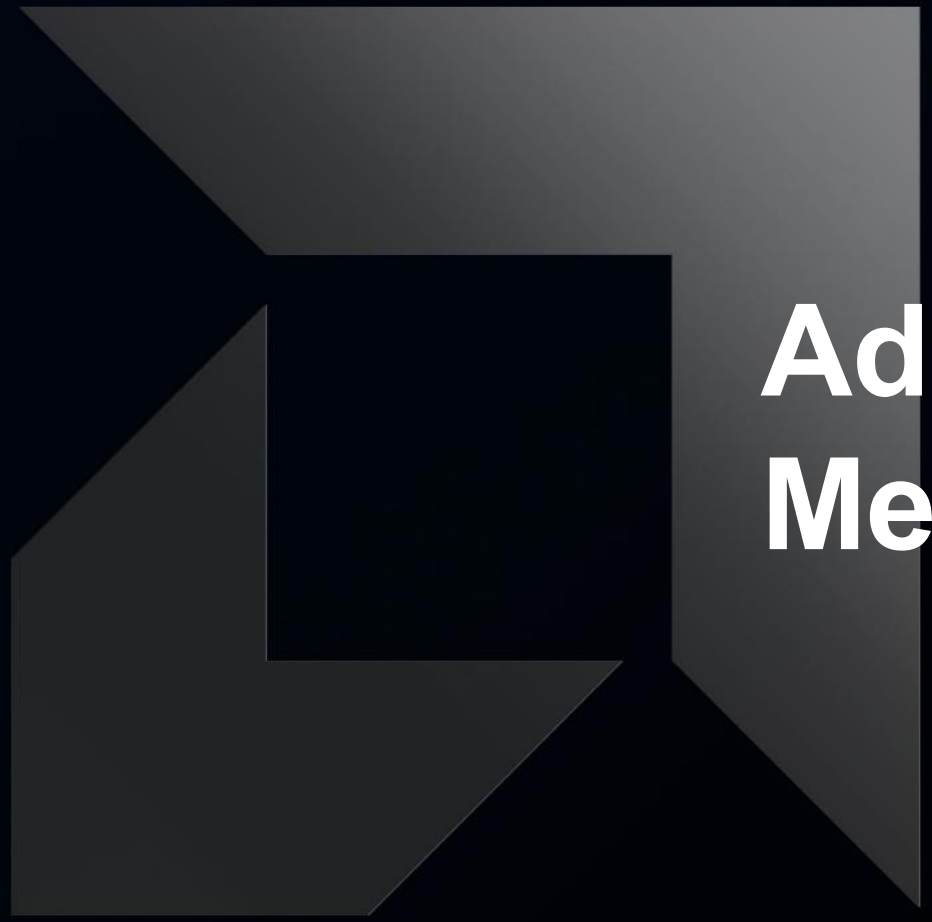

Addressing the Memory Evolution

AMD
together we advance_

The DRAM Market Remains Tight well into 2027

Channel Positioning

- DRAM supply remains constrained through 2027 due to stained AI-driven demand
- Limited supply expansion is driving structural pricing pressure (“memflation”)
- Memory inflation is a long-term shift, not a temporary market fluctuation

What This Means for the Channel

- Memory is becoming a primary driver of server TCO and refresh decisions
- Over-provisioned memory especially in 2P systems, increases cost without proportional benefit
- Partners need guidance to balance capacity, bandwidth, and platform efficiency

AMD Positioning & Guidance

- Lead with right-sized AMD EPYC™ Server CPU solutions that balance performance and memory efficiency
- Prioritize 1P platforms to deliver required performance with fewer DIMMs and lower system cost
- Optimize for memory utilization and bandwidth to reduce spend

Partner Outcome

- Lower customer TCO despite ongoing memory cost inflation
- Increased attach and differentiation through optimized platform design
- Shift customer conversations from “more memory” to “smarter memory”

Source: Gartner, Emerging Issue: Managing “Memflation” Through 2027 (27 Feb 2026)

Modernize Your Data Center to Stay Ahead

Act Now and Here is Why

- The Cost of Doing Nothing is rising
- Not modernizing introduces risk, higher costs, and missed efficiency gains

The Cost of Doing Nothing

- Rising operational risk from gaining infrastructure and security exposure
- Escalating support and warranty costs on legacy infrastructure
- Missed efficiency gains from modern, optimized platforms

Why Act Now

- Modern AMD EPYC™ Server CPU solutions reduce cost and improve performance per watt
- Right-sized infrastructure lowers memory, licensing, and system costs
- Improved security, reliability, and operational efficiency

AMD Advantage

- Strong TCO and rapid payback, even in today's memory cost environment
- AMD EPYC™ Server CPUs enable consolidation and right-sizing for immediate savings
- Proven path to modernize without sacrificing performance

Why Moving from 2P to 1P makes sense in a “Memflation” Environment

Customer Reality: Memory can now account for 50+% of Server BOMs

- 2P platforms scale memory by adding DIMMs to a second CPU. Cheap memory makes overprovisioning easy.
- Optimizing 2P platform performance with overprovisioned memory can now be cost prohibitive.
- 1P platforms allow ability to right-size memory footprint while optimizing platform performance.

Real-World Example: 2P → 1P

2P
- 16 or 32 DIMMs to populate all channels and optimize performance

1P AMD EPYC™ Server CPU
- 12 or 24 DIMMs to populate all channels and optimize performance

Fewer DIMMs = LESS MONEY to optimize performance without compromising compute or IO connectivity.

Typical Partner Outcomes

- Higher attach rates
- Faster refresh decisions
- Easier competitive displacement

Protect your customer’s budget by letting them buy fewer and lower-density DIMMs — with better performance. Lead with AMD EPYC™ Server CPU platforms and win both technically and financially.

1P AMD EPYC™ Server CPU 9355P vs 2P Intel Xeon® 5218 Exponential Benefits

Better Energy Efficiency

Save up to 56% on power vs. Intel based legacy systems[1]

Larger Space Savings

Up to 57% fewer servers, saving space and licensing cost[1]

Higher TCO Savings

Save up to 60%

TCO 5-Years savings vs. Intel based legacy systems[1]

More Performance Per Core

Up to 2.5x Faster

vs. Intel based legacy systems[1]

Lower Upfront Cost

Up to 68% Capex Savings

vs. Intel based legacy systems[1]

[1]1P AMD EPYC™ Server CPU 9355P 32C vs 2P Intel Xeon 5218 16Cdelivering 16740 integer score. See endnote 9xx5TCO-022

Compute Bound:
Refresh with Half the Memory/Core from 2P to 1P

Move to a 1P AMD EPYC™ Server CPU solution, less memory/core, and achieve payback in ~15.2 months

Legacy Intel
16GB memory/core

30 servers – **1,920** cores
286,105 KWh per year
13,380 SPECint

~Same Performance

Half Memory/Core

~49% Less Power

New AMD
8GB memory/core

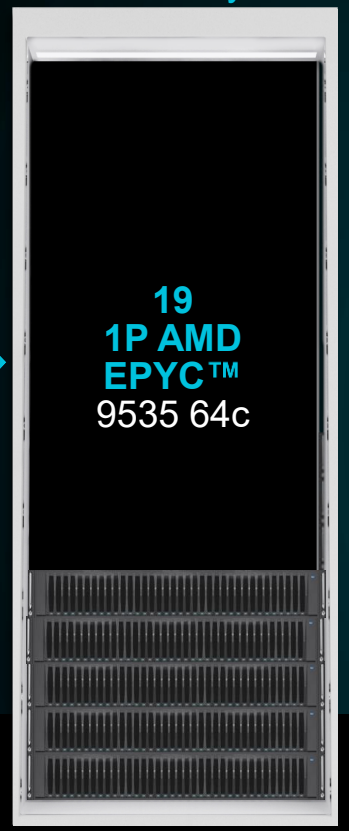

19 servers – **1,216** cores
147,246 KWh per year
13,380 SPECint

Systems & SW

~32%
Lower

5 Yr TCO

With Licensing Costs:
Payback: ~15.2 months

As of 4/7/2026. Analysis based on the AMD EPYC™ Server Greenhouse Gas Emission TCO Estimation Tool - version 1.55

Memory Bound:
Refresh with the Same Memory/Core (64C)

Maintain memory/core and achieve payback in ~45 months

Legacy Intel
16GB/core memory

30 servers – **1,920** cores
286,105 KWh per year
13,380 SPECint

~Same Performance

Same Memory/Core

~53% Less Power

New AMD
16GB/core memory

17 servers – **1,088** cores
134,784 KWh per year
13,380 SPECint

Systems & SW

~30%
Lower

5 Yr TCO

With Licensing Costs:
Payback: ~45 months

As of 4/7/2026. Analysis based on the AMD EPYC™ Server Greenhouse Gas Emission TCO Estimation Tool - version 1.55

Stop Over-Provisioning Memory — Right-Size for the Workload

Most environments are over-provisioned on memory—right-sizing protects performance and reduces cost.”.

Compute-Bound Workloads

- Memory is typically over-provisioned relative to compute needs
- Significant memory reduction possible with minimal performance impact
- Opportunity to lower cost while maintaining throughput

Balanced/Mixed Workloads

- Moderate memory optimization delivers strong cost performance balance
- Right-sizing aligns infrastructure to actual workload demand
- Enables efficient scaling without over-investment

Memory Bound Workloads

- Memory remains a key performance driver
- Selective optimization can reduce cost
- Right-sized configuration to avoid unnecessary over-provisioning

Right-Size Your Memory for your Workload with AMD EPYC™ Server CPUs

Compute Bound Workloads

50% Memory Capacity Reduction

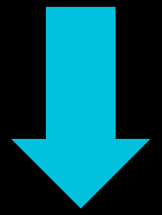

2% Performance Reduction

Tested Workloads: SIR2017, SpecJBB 2015, FFMpeg, NGINX, REDIS, MySQL TPC-C

Memory Capacity Bound Workloads

25% Memory Capacity Reduction

5% Performance Reduction

50% Memory Capacity Reduction

25% Performance Reduction

Tested workload: SQL Server OLAP

Compute bound = Right size your memory, **don't over provision**

Observed performance impact is a geomean across the tested workloads

Optimize Memory for the Workload Capacity or Bandwidth, Not Both by Default

Memory-Intensive / Capacity-Driven (Tiering-Friendly)

Best-Fit Workloads

Relational Databases (Oracle, SQL Server)
Capacity-driven instances, reporting, consolidation

In-Memory Databases
(Redis, MongoDB, TPC-H) Lg Memory Footprint Databases (Hybrid / Tier-Aware)(Analytics, read-heavy, warm/cold datasets)

First-Party Enterprise Workloads
(Internal cloud infra, business apps,)

Lead with capacity-optimized EPYC for dataset-driven, tier-friendly workloads.

AMD Guidance
- Prioritize right-sized memory capacity, not peak bandwidth
- Leverage fewer sockets (1P where possible) to reduce DIMM count and platform cost (Fewer NUMA domains = simpler tuning).
- Use memory tiering to lower TCO without disruption applications

Memory-Bandwidth Bound
(when bandwidth truly matters)

Best-Fit Workloads

Cloud General Purpose
(VM density, burst workloads)

Automotive / Manufacturing
(ANSYS Mechanical, LS-DYNA)

Electronic Design Automation
(RTL simulation, IR/EM analysis)

Optimize for bandwidth when performance depends on data movement.

AMD Guidance
- Size memory to bandwidth needs, not historical capacity assumptions
- Favor balanced DIMM configurations that sustain throughput at scale
- Position AMD EPYC™ Server CPU for higher VM density and efficient platform utilization

Capacity drives cost. Bandwidth drives performance. EPYC lets you optimize for the one that matters.


AMD
together we advance_data centers