



5 REASONS TO SWITCH FROM INTEL® XEON® TO AMD EPYC™ PROCESSORS

Savvy enterprise IT leaders know that if you aren't challenging the status quo, you're likely falling behind.

Making informed, strategic decisions about CPU and GPU providers helps you take data centers to the next level. With AMD you can plan and execute with confidence, streamline costs, and make future-ready investments.

1

SHRINK YOUR DATA CENTER FOOTPRINT UP TO 7X: CHOOSE TO MODERNIZE

The hardware you use makes all the difference in combating costs and meeting strict project deadlines and budget requirements. A single AMD EPYC™ 9005 CPU-based server can do the work of more than seven 2019-era Intel® Xeon® Platinum servers,¹ which can lower your energy consumption and shrink your data center footprint, freeing up space for new server infrastructure to meet growing AI workload demands.



CONSUME UP TO 69% LESS POWER

Replacing 1,000 2P Intel Xeon Platinum 8280 CPU-based servers with 127 2P EPYC 9965 CPU-based servers to deliver 391,000 units of integer performance uses up to 69% less electricity, saving up to \$3.4M in energy bills over five years.²

2

DON'T MISS OUT ON AI'S POTENTIAL: CHOOSE AN AI LEADER

Confidently run AI workloads today with available, cost-efficient, and high-performance offerings from a leading supplier. AMD has focused its hardware innovation on AI, delivering optimized solutions such as AMD Instinct™ accelerators combined with EPYC CPUs to help you get the most out of your AI investment.



HANDLE UP TO 700K MORE TOKENS PER SECOND

A 1,000 node AI cluster using 8x AMD Instinct MI300X accelerators paired with EPYC 9575F CPUs delivers 700,000 more tokens per second compared to the same cluster using Intel Xeon Platinum 8592+ CPUs when running Llama 3.1-70B (128/2048 I/O tokens, FP8).³

3

AVOID PROPRIETARY HARDWARE: CHOOSE OPEN SOLUTIONS

With the pace of technology evolution, especially for AI, being tied down to a single supplier carries the risk of overspending and missing out on upgrade cycles. Protect your investments with AMD, a supplier that innovates and provides flexibility with standards-based hardware—including for Ultra Accelerator Link, JEDEC memory, and Ultra Ethernet—in open software environments.



DID YOU READ THE FINE PRINT?

Memory and networking are critical factors in hardware decision-making. Make sure the solutions you select come with JEDEC-standard memory and standard Ethernet instead of expensive proprietary alternatives.

4

COUNT ON THE RIGHT SOLUTION: CHOOSE A COMPREHENSIVE PORTFOLIO

No matter what your unique workloads require, AMD can support your needs with an end-to-end portfolio and open ecosystem. Our diverse and mature portfolio offers versatile data center hardware, including all-purpose CPUs for on-premises, cloud, and AI demands; a premier line of GPUs for AI; and robust, interoperable networking solutions.



EARNING 35% OF THE DATA CENTER MARKET IS JUST THE START⁴

Enterprise customers are seeing the advantages of AMD and adopting EPYC CPU-based solutions to power services used by billions every day.

5

DON'T LET DEPLOYMENT PLANS GET DERAILED: CHOOSE A SUPPLIER WITH A RELIABLE TRACK RECORD

Tired of constantly changing your IT plans because of supplier constraints and lack of execution? Join the growing number of industry leaders around the world that have mitigated the risk of disruption by choosing AMD EPYC servers.



EXECUTION YOU CAN COUNT ON

Depend on a server platform supplier that delivers on its promises—AMD delivered five generations of leadership EPYC processors to market on time, advancing data center innovation worldwide.

THE SMARTEST CHOICE? **AMD.**

The decisions you make as you design and manage data centers set you up to succeed—or to struggle. Join the ranks of other IT professionals who have learned how to use AMD CPUs and GPUs to improve efficiency, performance, and costs. Work with your server manufacturer or solution builder to match AMD products to your current and future workloads.

Remember: You decide what goes inside your servers. [Explore AMD EPYC CPUs.](#)

1. This scenario contains many assumptions and estimates and, while based on AMD internal research and best approximations, should be considered an example for information purposes only, and not used as a basis for decision making over actual testing. The AMD Server & Greenhouse Gas Emissions TCO (total cost of ownership) Estimator Tool - version 1.3, compares the selected AMD EPYC™ and Intel® Xeon® CPU based server solutions required to deliver a TOTAL PERFORMANCE of 391000 units of SPECrate[®]2017_int_base performance as of November 21, 2024. This estimation compares upgrading from a legacy 2P Intel Xeon 28 core Platinum 8280 based server with a score of 391 (<https://spec.org/cpu2017/results/res2024q4/cpu2017-20200915-23984.pdf>) versus 2P EPYC 9965 (192C) powered server with a score of 3100 (<https://spec.org/cpu2017/results/res2024q4/cpu2017-20241004-44979.pdf>). Environmental impact estimates made leveraging this data, using the Country / Region specific electricity factors from Country Specific Electricity Factors - 2024, and the United States Environmental Protection Agency Greenhouse Gas Equivalencies Calculator. For additional details, see <https://www.amd.com/en/legal/claims/epyc.html#q=9xx5TCO-005>. (9xx5TCO-005)
2. Ibid.
3. As of 10/10/2024; this scenario contains several assumptions and estimates and, while based on AMD internal research and best approximations, should be considered an example for information purposes only, and not used as a basis for decision making over actual testing. Referencing 9XX5-056A: "2P AMD EPYC 9575F powered server and 8x AMD Instinct MI300X GPUs running Llama3.1-70B select inference workloads at FP8 precision vs 2P Intel Xeon Platinum 8592+ powered server and 8x AMD Instinct MI300X GPUs has ~8% overall throughput increase across select inference use cases" and 8763.52 tokens/s (9575F) versus 8,048.48 tokens/s (8592+) at 128 input / 2048 output tokens, 500 prompts for 1.089x the tokens/s or 715.04 more tokens/s. 1 Node = 2 CPUs and 8 GPUs. Assuming a 1000 node cluster, 1000 * 715.04 = 715,040 tokens/s For ~700,000 more tokens/s. Results may vary due to factors including system configurations, software versions and BIOS settings. (9xx5-087)
4. Mercury research sell-in revenue shipment estimates, 2024 Q4.

© 2025 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, EPYC, AMD Instinct, and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and other countries. SPEC[®], SPEC CPU[®], and SPECrate[®] are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. Other product names used in this publication are for identification purposes only and may be trademarks of their respective owners.