together we advance_

DESIGN WITH AMD
SEE HOW YOUR NEXT DATA CENTER CAN *PAY FOR ITSELF*

Deciding what goes inside your next server is critical.
Choosing the right **AMD CPUs and GPUs** for your workloads can increase performance and drastically reduce operating costs.

Choose the right hardware and save over the life of your next data center

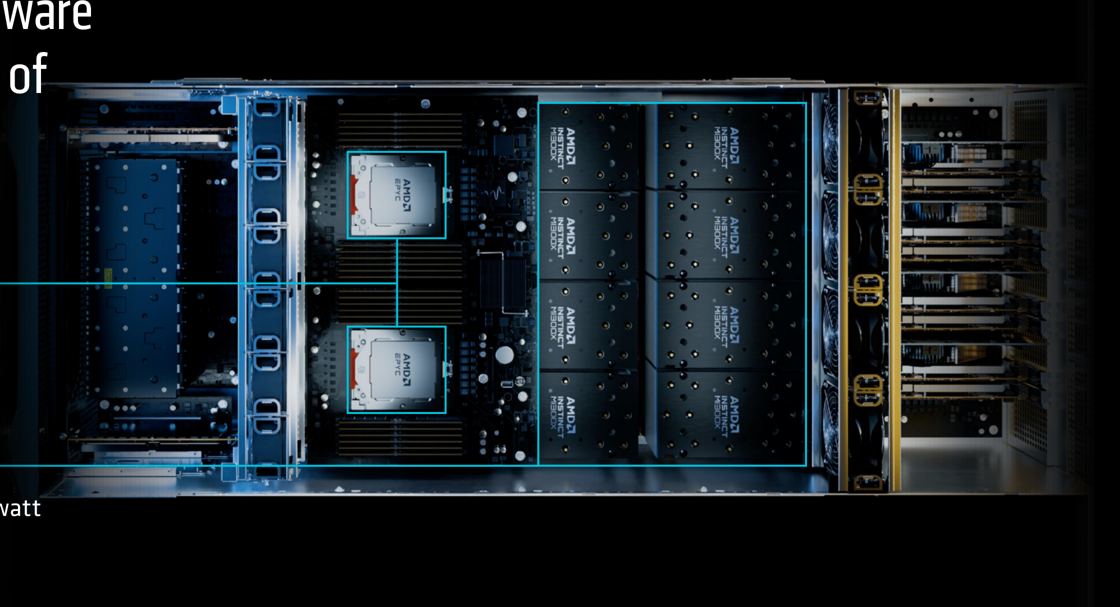

CPUs
CPUs drive hardware efficiency gains for enterprise workloads[1]

GPUs
GPUs can improve performance-per-watt exponentially for AI workloads

YOU DECIDE:
Yesterday's tech or advanced performance

LEGACY x86 SERVERS	*AMD EPYC™ CPU-BASED SERVERS*
Yesterday's performance, years behind schedule	*Industry-leading performance with headroom for the future*
Bottlenecked AI accelerator performance	*Higher throughput from GPU platforms*
Ballooning virtual machine costs	*Lower software licensing costs*
Power-hungry servers	*Reduced energy bills for the life of the data center*
Dead ends are built in	*Open standards for flexibility*

UP TO **2.7X** SPEC® CPU PERFORMANCE[2]
5th Generation AMD EPYC™ 9965 CPUs outperform 5th Gen Intel® Xeon® Platinum 8592+ CPUs.

UP TO **20%** MORE AI INFERENCE THROUGHPUT[3]
8x GPU platforms process up to 20% more tokens per second with 2P AMD EPYC 9575F CPU hosts versus 2P Intel Xeon Platinum 8592+ hosts.

UP TO **$44M** IN SOFTWARE SAVINGS OVER FIVE YEARS[4]
Replacing 1,000 2P Intel Xeon Platinum 8280 servers with 240 2P AMD EPYC 9555 servers can save up to $44 million in virtual machine software costs over five years.

UP TO **$3.4M** IN ENERGY SAVINGS OVER FIVE YEARS[5]
Replacing 1,000 2P Intel Xeon Platinum 8280 servers with 127 2P AMD EPYC 9965 servers delivers similar performance with far less electricity.

PROTECT YOUR INVESTMENT
Keep your options open at build time and over the life of your data center.

HOW DATA CENTERS POWERED BY AMD EPYC *HELP PAY FOR THEMSELVES*

One new AMD server can do the work of seven older servers[6]

1,000 2P INTEL XEON PLATINUM 8280 SERVERS = **127** 5TH GENERATION 2P AMD EPYC 9965 SERVERS

UP TO **67%** LOWER TCO OVER FIVE YEARS[7] + UP TO **68%** LESS ENERGY CONSUMPTION[8] + UP TO **87%** FEWER SERVERS[9]

Choose to upgrade to AMD EPYC CPUs and save millions in energy and software costs over five years[10]

HARDWARE SAVINGS: $871,417
POWER SAVINGS: $219,835
VIRTUALIZATION SOFTWARE LICENSING SAVINGS: $6,944,000
TOTAL FIVE-YEAR SAVINGS: $8,035,252

Initial cost for AMD EPYC CPU-powered servers: $2,218,671
FIVE-YEAR ROI: *$5,816,581 – 2.6X INITIAL CAPEX*

Estimated costs compare 69 2P AMD EPYC 9575F CPU-powered servers to 100 Intel Xeon 8592 CPU-powered servers delivering equivalent total performance.

REMEMBER:
You decide what goes inside

Never settle. Dig into the hardware details and work with your server manufacturer to choose the right AMD EPYC CPU-based server solutions for your workloads.

Explore AMD EPYC PROCESSORS

1. SPECrate®2017_int_base comparison based on published scores from www.spec.org as of 10/10/2024. 2P AMD EPYC 9965 (3100 SPECrate®2017_int_base, 384 Total Cores, 500W TDP, $14,813 CPU $), 6.200 SPECrate®2017_int_base/CPU W, 0.200 SPECrate®2017_int_base/CPU $ (https://www.spec.org/cpu2017/results/res2024q4/cpu2017-20241004-44979.html); 2P AMD EPYC 9755 (2720 SPECrate®2017_int_base, 256 Total Cores, 500W TDP, $12,984 CPU $), 5.440 SPECrate®2017_int_base/CPU W, 0.209 SPECrate®2017_in t_base/CPU $ (https://www.spec.org/cpu2017/results/res2024q4/cpu2017-20240923-44824.html); 2P AMD EPYC 9754 (1950 SPECrate®2017_int_base, 256 Total Cores, 360W TDP, $11,900 CPU $), 5.417 SPECrate®2017_int_base/CPU W, 0.164 SPECrate®2017_int_base/CPU $ (https://www.spec.org/cpu2017/results/res2023q2/cpu2017-20230522-36617.html); 2P AMD EPYC 9654 (1810 SPECrate®2017_int_base, 192 Total Cores, 360W TDP, $11,805 CPU $), 5.028 SPECrate®2017_int_base/CPU W, 0.153 SPECrate®2017_int_base/CPU $ (https://www.spec.org/cpu2017/results/res2024q1/cpu2017-20240129-40896.html); 2P Intel Xeon Platinum 8592+ (1130 SPECrate®2017_int_base, 128 Total Cores, 350W TDP, $11,600 CPU $) 3.229 SPECrate®2017_int_base/CPU W, 0.097 SPECrate®2017_int _base/CPU $ (https://spec.org/cpu2017/results/res2023q4/cpu2017-20231127-40064.html); 2P Intel Xeon 6780E (1410 SPECrate®2017_int_base, 288 Total Cores, 330W TDP, $11,350 CPU $) 4.273 SPECrate® 2017_int_base/CPU W, 0.124 SPECrate®2017_int_base/CPU $ (https://spec.org/cpu2017/results/res2024q3/cpu2017-20240811-44406.html). SPEC®, SPEC CPU®, and SPECrate® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. Intel CPU TDP at https://ark.intel.com/. (9xx5-002E)
2. Ibid.
3. Llama3.1-70B inference throughput results based on AMD internal testing as of 09/01/2024. Llama3.1-70B configurations: TensorRT-LLM 0.9.0, nvidia/cuda 12.5.0-devel-ubuntu22.04, FP8, Input/Output token configurations (use cases): [BS=1024 I/O=128/128, BS=1024 I/O=128/2048, BS=96 I/O=2048/128, BS=64 I/O=2048/2048]. Results in tokens/second. 2P AMD EPYC 9575F (128 Total Cores) with 8x NVIDIA H100 80GB HBM3, 1.5TB 24x64GB DDR5-6000, 1.0 Gbps 3TB Micron_9300_MTFDHAL3T8TDP NVMe®, BIOS T20240805173113 (Determinism=Power,SR-IOV=On), Ubuntu 22.04.3 LTS, kernel=5.15.0-117-generic (mitigations=off, cpupower frequency-set -g performance, cpupower idle-set -d 2, echo 3> /proc/syss/vm/drop_caches), 2P Intel Xeon Platinum 8592+ (128 Total Cores) with 8x NVIDIA H100 80GB HBM3, 1TB 16x64GB DDR5-5600, 3.2TB Dell Ent NVMe® PM1735a MU, Ubuntu 22.04.3 LTS, kernel-5.15.0-118-generic, (processor.max_cstate=1, intel_idle.max_cstate=0 mitigations=off, cpupower frequency-set -g performance), BIOS 2.1, (Maximum performance, SR-IOV=On), I/O Tokens Batch Size EMR Turin Relative Difference 128/128 1024 814.678 1101.966 1.353 287.288 128/2048 1024 2120.664 2331.776 1.1 211.112 2048/128 96 114.954 146.187 1.272 31.233 2048/2048 64 333.325 354.208 1.063 20.833 For average throughput increase of 1.197x. When scaling to a 1000 node cluster (1 node = 2 CPUs and 8 GPUs) comparing the AMD EPYC 9575F system and Intel Xeon 8592+ system: 128/128 achieves 287,288 more tokens/s, 128/2048 achieves 211,112 more tokens/s, 2048/128 achieves 31,233 more tokens/s, 2048/2028 achieves 20,833 more tokens/s. Results may vary due to factors including system configurations, software versions and BIOS settings. (9xx5-014A)
4. 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/res2020q3/cpu2017-20200915-23984.pdf) versus 2P EPYC 9555 (128C) powered server with a score of 1630 (https://spec.org/cpu2017/results/res2024q4/cpu2017-20241104-45226.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-006. (9xx5TCO-006)
5. Ibid.
6. 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 1.12, compares the selected AMD EPYC™and Intel® Xeon® CPU based server solutions required to deliver a TOTAL_PERFORMANCE of 391000 units of SPECrate2017_int_base performance as of October 10, 2024. This estimation compares a legacy 2P Intel Xeon 28 core Platinum_8280 based server with a score of 391 versus 2P EPYC 9965 (192C) powered server with a score of 3000 (https://www.spec.org/cpu2017/results/res2024q4/cpu2017-20240923-44837.pdf) along with a comparison upgrade to a 2P Intel Xeon Platinum 8592+ (64C) based server with a score of 1130 (https://spec.org/cpu2017/results/res2024q3/cpu2017-20240701-43948.pdf). Actual SPECrate®2017_int_base score for 2P EPYC 9965 will vary based on OEM publications. Environmental impact estimates made leveraging this data, using the Country/ Region specific electricity factors from the 2024 International Country Specific Electricity Factors 10 –July 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-002a. (9xx5TCO-002A)
7. Ibid.
8. Ibid.
9. Ibid.
10. 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 ~113000 units of SPECrate®2017_int_base performance as of January 6, 2025. This estimation compares a 2P Intel Xeon 64 core Platinum_8592+ based server with a SPECrate2017_int_base score of 1130, https://spec.org/cpu2017/results/res2023q4/cpu2017-20231127-40064.pdf; versus 2P EPYC 9575F (64C) powered server with a score of 1640 https://spec.org/cpu2017/results/res2024q4/cpu2017-20241104-45229.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-010. (9xx5TCO-010)

together we advance_