WORLD'S BEST CPU FOR THE AGENTIC AI ERA: THE NUMBERS BEHIND THE LEADERSHIP

AMD EPYC™ 9006 SP7 Server CPUs delivers leadership compute density and performance across the CPU roles that scale agents, maximize accelerator throughput, and power enterprise AI workloads.[1]

AMD
together we advance_

THE CPU FOUNDATION BEHIND AGENTIC AI

Running agents is only part of the picture. Infrastructure also needs to keep accelerators utilized and support the databases, applications, and services agents rely on.

Explore the benchmarks across three essential CPU roles:
AI Host Node, Agent Sandbox, and **General-Purpose CPUs.**

RUN MORE AGENTS.
PER WATT. PER DOLLAR. PER RACK.

2.9x the threads per MW of system power versus **NVIDIA Vera**[2], and 2.0x the SPECrate®2026_int_base performance per MW of system power

GET MORE FROM
GPU ACCELERATORS

1.26x the per-core **SPECrate®2026_int_base performance over NVIDIA Vera**[3], enabled by frequencies up to 5.3GHz

MORE JOBS DONE.
FASTER INSIGHTS. FEWER SERVERS.

2.8x the SPECrate®2026_int_base throughput of NVIDIA Vera[4]

1 AI HOST NODE CPU
KEEP ACCELERATORS UTILIZED

AI Host Node CPUs help keep accelerators supplied with data and running at full potential. These benchmarks highlight throughput and per-core performance that maximize accelerator utilization and token generation performance for agentic AI workloads.

PLATFORM ADVANTAGE WITH FASTEST MEMORY BANDWIDTH

Up to 1.6 TB/s per-socket memory bandwidth with first to market support for JEDEC-standard MRDIMM-12800, supporting large-context agentic AI reasoning[5]

PER-CORE PERFORMANCE

1.26X Per-core performance over NVIDIA Vera[6], enabled by frequencies up to 5.3GHz

2 AGENT SANDBOX CPU
SCALE MORE AGENTS PER RACK

Agent Sandbox CPUs maximize agent capacity by scaling the number of agents that can run within power, cost and rack constraints. These numbers focus on concurrency, density, and power efficiency at rack scale.

RACK-SCALE POWER EFFICIENCY

2.9X the threads per MW of system power versus NVIDIA Vera[7]

CONCURRENT AGENT CAPACITY

1.6X More concurrent agents per Helion rack than NVIDIA Vera-based systems[8]

3 GENERAL-PURPOSE CPU
RUN THE DIVERSE WORKLOADS AGENTS DEPEND ON

General-Purpose CPUs provide the performance for virtualization, databases, storage, CPU inference, agentic orchestration, and complex HPC workloads.

GENERAL-PURPOSE SPECint LEADERSHIP

2.3X throughput SPECrate® throughput vs Intel Xeon 6980P.[9]

COMPLEX HPC WORKLOADS

Up to 3.2X the performance of Intel Xeon 6980P across life sciences, quantum chemistry, and weather forecasting.[10]

AGENTIC AI PIPELINE PERFORMANCE

2.74X the performance of Intel® Xeon® 6980P[11] Geometric mean across five CPU centric agentic AI pipeline tiers

AGENTIC AI IS A PIPELINE, NOT A SINGLE WORKLOAD.

Take a look at the AMD EPYC™ Server CPUs and the workload-by-workload data behind these numbers.

1.3x *THE PER-CORE PERFORMANCE*

Per-core leadership vs NVIDIA Vera
Benchmark: SPECrate®2017_int_base (estimated)[12]

Throughput leadership vs NVIDIA Vera
Benchmark: SPECrate®2026_int_base[4]

2.8x *HIGHER THROUGHPUT*

2.2x *THE PERFORMANCE OF 2P INTEL XEON 6980P*

Dual-socket performance leadership vs Intel
Benchmark context: 2P estimated score comparison[13]

Enterprise and cloud server workload uplift
Benchmark context: fixed-frequency IPC across 27 workloads[14]

~27% *IPC UPLIFT*

RACK-LEVEL COMPARISON

Performance Normalized vs. NVIDIA Vera



READY TO SIZE YOUR INFRASTRUCTURE FOR THE AGENTIC ERA?

Explore the workload-by-workload data behind these numbers and find the right AMD EPYC 9006 Series Server CPU for every role in your agentic AI workflow.

See the full benchmark

[1] 9xx6-005C: as of 11/5/2026, AMD EPYC CPUs can support up to 1.33x the vCPUs and threads vs 5th Gen AMD EPYC 9965. 6th Gen AMD EPYC 9996 has the highest thread density (thread count) and the highest multithreaded core density of server CPUs in production. vCPU support based on thread count, comparing Top of Stack 6th Gen EPYC 9996 CPU with 256 cores (512 threads) and 5th Gen EPYC 9965 with 192 cores (384 threads). In comparison, the highest thread count Intel® Xeon® processor is the Intel Xeon 6990E at 288 cores (288 threads). The highest multithreaded core count is the Intel Xeon 6980P at 128C (256 threads), and the highest thread count Nvidia CPU is the Nvidia Vera with 88 cores (176 threads). For reference, 6th Gen Intel Xeon CPUs support up to ~0.8x, ARM AGI with 136 cores (136 threads, 1 thread per core) is expected to support up to ~0.35x, and Nvidia Vera is expected to support ~0.5x the thread count of 5th Gen AMD EPYC CPUs. Intel specifications from ark.intel.com. ARM specifications from https://www.arm.com/products/cloud-datacenter/arm-agi-cpu#Specifications. Nvidia specifications from https://developer.nvidia.com/blog/nvidia-vera-cpu-delivers-high-performance-bandwidth-and-efficiency-for-ai-factories/.

[2] 9xx6-063A: Comparison based on AMD internal testing and estimates and published AMD and NVIDIA specifications as of 11/05/2026. AMD EPYC scores are estimates based on AMD internal testing. System power per CPU = CPU power / 85% voltage regulator efficiency + memory power at 65% load + NIC + storage + BMC + other. 2P 6th Gen AMD EPYC 256C at 400W: 32x64GB DDR5-8000 RDIMM, GCC 16.1, estimated SPECrate®2026_int_base 1810, 1,417 W system power per 2P server (708 W per CPU). NVIDIA Vera: 88C, 450W TDP, ETL256 rack (8-CPU system), LPDDR5X-9600, GCC 15.2, estimated SPECrate®2026_int_base 925 (2P), 1,417 W system power per 2 CPUs (709 W per CPU). Within 1 MW of system power: 706 AMD EPYC 2P servers × 1,024 threads = 722,768 threads vs 706 Vera 2-CPU equivalents × 352 threads = 248,403 threads, 2.9x. Performance within 1 MW: 706 × 1810 = 1,277,550 vs 706 × 925 = 652,764 SPECrate®2026_int_base, 2.0x. Cores per rack: AMD Helion rack 32,768 cores, https://www.amd.com/en/products/rackscale-solutions/helion.html#configurations, vs NVIDIA Vera ETL256 rack 22,528 cores (256 CPUs × 88 cores), https://www.nvidia.com/en-us/data-center/products/vera-rack/. 32,768 / 22,528 = 1.45x. Agents per rack: AMD Helion rack 65,536 threads vs NVIDIA Vera ETL256 rack 45,056 threads, 2 threads per core (SMT) on both, ~1.45x. Agent counts are estimates derived from available CPU thread resources used as a proxy under a consistent theoretical workload. Actual agent capacity and throughput will vary based on workload, model, memory, software, orchestration, and system configuration. Nvidia Vera performance from: https://nvdam.widen.net/s/nmw5vblpqd/nvidia_vera_cpu_architecture_whitepaper?nvid=nv-tblg-543584. SPEC®, SPEC CPU®, and SPECrate® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. Results may vary.

[3] 9xx6-049D: SPECrate®2026_int_base per-core comparison based on the published 2P 6th Gen AMD EPYC 9686F result, the highest published 96-core Intel Xeon SPECrate®2026_int_base result and the highest published 2P Intel Xeon 6980P result from www.spec.org as of 11/05/2026, and NVIDIA published estimates for NVIDIA Vera. Per-core results use total processor cores (per socket x sockets). 2P AMD EPYC 9686F (GCC 16.1) delivers 1.26x the per-core performance of 2P NVIDIA Vera (GCC 15.2), 1.73x the per-core performance of 2P Intel Xeon 6952P (Intel oneAPI 2025.3) and 1.51x the per-core performance of 2P Intel Xeon 6980P (Intel oneAPI 2026.0). Configurations: 2P AMD EPYC 9686F: 96 cores per socket (192 cores total), SMT on, 32x 64GB DDR5-8000 RDIMM, 500W Default CPU Power, power determinism mode, GCC 16.1. 384 copies, https://www.spec.org/cpu2026/results/res2026q4/cpu2026-20261020-00766.html 2P NVIDIA Vera: 88 cores per socket (176 cores total), GCC 15.2, NVIDIA published estimate, https://nvdam.widen.net/s/nmw5vblpqd/nvidia_vera_cpu_architecture_whitepaper?nvid=nv-tblg-543584 2P Intel Xeon 6952P: Dell PowerEdge R770AP, 96 cores per socket (192 cores total), SMT on, 400W TDP, 768GB (24x 32GB DDR5-6400), Intel oneAPI 2025.3, 384 copies, published May-2026, https://www.spec.org/cpu2026/results/res2026q2/cpu2026-20261020-00069.html 2P Intel Xeon 6980P: 128 cores per socket (256 cores total), 500W TDP, 24x 64GB MRDIMM-8800, Intel oneAPI 2026.0, published Jun-2026, https://spec.org/cpu2026/results/res2026q2/cpu2026-20260519-00255.html Starting with the 6th Gen AMD EPYC™ server processor family, AMD uses Default CPU Power to describe processor power consumption, succeeding AMD's historical TDP reference. SPEC®, SPEC CPU® and SPECrate® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. Results may vary due to factors including system configurations, software versions and BIOS settings. 9xx6-030A: Peak frequency comparison based on published specifications as of October 13, 2026. 6th Gen AMD EPYC 9686F (96C, high-frequency): 5.3 GHz max boost frequency. Source: amd.com. Intel Xeon 6972P (96C): 3.9 GHz max turbo frequency. Source: https://ark.intel.com. 5th Gen AMD EPYC 9575F (64C, high-frequency): 5.0 GHz max boost frequency. Source: amd.com. Intel Xeon 6960P (72C): 3.9 GHz max turbo frequency. Source: https://ark.intel.com. As of October 13, 2026, no other publicly available server processor with 96 or more cores offers a 5.3 GHz max frequency (Intel Xeon 6972P 96C: 3.9 GHz; NVIDIA Vera 88C: up to 3.2 GHz; Arm AGI 136C: up to 3.7 GHz; Intel Xeon 6776P with PCT: 4.6 GHz at 64C). Max boost for AMD EPYC processors is the maximum frequency achievable by any single core on the processor under normal operating conditions for server systems. Intel and Xeon are trademarks of Intel Corporation or its subsidiaries.

[4] 9xx6-107: SPECrate®2026_int_base comparison based on the highest published 2P AMD EPYC 9996 score and the published 2P AMD EPYC 9965 AOCC 5.1.0 score from www.spec.org as of 11/05/2026, and the NVIDIA published estimate for 2P NVIDIA Vera. 2P AMD EPYC 9996, 512C, AOCC 6.0, SPECrate®2026_int_base 2560, https://spec.org/cpu2026/results/res2026q4/cpu2026-20261014-00694.html 2P AMD EPYC 9965, 384C, AOCC 5.1.0, SPECrate®2026_int_base 1270, https://spec.org/cpu2026/results/res2026q2/cpu2026-20260210-00037.html 2P NVIDIA Vera, 176C, GCC 15.2, NVIDIA published estimate SPECrate®2026_int_base 925, https://nvdam.widen.net/s/nmw5vblpqd/nvidia_vera_cpu_architecture_whitepaper?nvid=nv-tblg-543584 SPEC®, SPEC CPU® and SPECrate® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. Results may vary due to factors including system configurations, software versions and BIOS settings.

[5] 9xx6-006A: 6th Gen EPYC CPUs can support up to 2.6X memory bandwidth (up to 16 channels of 12,800MT/s DDR5 for 1.6TB/s of memory bandwidth/socket compared to 614GB/s memory bandwidth/socket of 5th Gen EPYC CPUs (9005) with 12 channels of 6,400MT/s DDR5). 6th Gen EPYC CPUs can support up to 2X the IO bandwidth (PCIe Gen 6 supported at 64GT/s per lane. EPYC 9005 platforms support PCIe Gen 5 at 32GT/s per lane). Competitor memory bandwidth/socket: Vera: ~1.2 TB/s ARM AGI = 12 * 8 * 8800 = 844800 = ~0.8 TB/s 6980P = 12 * 8 * 8800 = 844800 = ~0.8 TB/s Intel specifications from ark.intel.com. ARM specifications from https://www.arm.com/products/cloud-datacenter/arm-agi-cpu#Specifications. Nvidia specifications from https://developer.nvidia.com/blog/nvidia-vera-cpu-delivers-high-performance-bandwidth-and-efficiency-for-ai-factories/. ARM AGI specifications from: https://www.arm.com/static/az/pdf/product-brief/arm-agi-cpu-product-brief.pdf

[6] 9xx6-049D: SPECrate®2026_int_base per-core comparison based on the published 2P 6th Gen AMD EPYC 9686F result, the highest published 96-core Intel Xeon SPECrate®2026_int_base result and the highest published 2P Intel Xeon 6980P result from www.spec.org as of 11/05/2026, and NVIDIA published estimates for NVIDIA Vera. Per-core results use total processor cores (per socket x sockets). 2P AMD EPYC 9686F (GCC 16.1) delivers 1.26x the per-core performance of 2P NVIDIA Vera (GCC 15.2), 1.73x the per-core performance of 2P Intel Xeon 6952P (Intel oneAPI 2025.3) and 1.51x the per-core performance of 2P Intel Xeon 6980P (Intel oneAPI 2026.0). Configurations: 2P AMD EPYC 9686F: 96 cores per socket (192 cores total), SMT on, 32x 64GB DDR5-8000 RDIMM, 500W Default CPU Power, power determinism mode, GCC 16.1. 384 copies, https://www.spec.org/cpu2026/results/res2026q4/cpu2026-20261020-00766.html 2P NVIDIA Vera: 88 cores per socket (176 cores total), GCC 15.2, NVIDIA published estimate, https://nvdam.widen.net/s/nmw5vblpqd/nvidia_vera_cpu_architecture_whitepaper?nvid=nv-tblg-543584 2P Intel Xeon 6952P: Dell PowerEdge R770AP, 96 cores per socket (192 cores total), SMT on, 400W TDP, 768GB (24x 32GB DDR5-6400), Intel oneAPI 2025.3, 384 copies, published May-2026, https://www.spec.org/cpu2026/results/res2026q2/cpu2026-20261020-00069.html 2P Intel Xeon 6980P: 128 cores per socket (256 cores total), 500W TDP, 24x 64GB MRDIMM-8800, Intel oneAPI 2026.0, published Jun-2026, https://spec.org/cpu2026/results/res2026q2/cpu2026-20260519-00255.html Starting with the 6th Gen AMD EPYC™ server processor family, AMD uses Default CPU Power to describe processor power consumption, succeeding AMD's historical TDP reference. SPEC®, SPEC CPU® and SPECrate® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. Results may vary due to factors including system configurations, software versions and BIOS settings. 9xx6-030A: Peak frequency comparison based on published specifications as of October 13, 2026. 6th Gen AMD EPYC 9686F (96C, high-frequency): 5.3 GHz max boost frequency. Source: amd.com. Intel Xeon 6972P (96C): 3.9 GHz max turbo frequency. Source: https://ark.intel.com. 5th Gen AMD EPYC 9575F (64C, high-frequency): 5.0 GHz max boost frequency. Source: amd.com. Intel Xeon 6960P (72C): 3.9 GHz max turbo frequency. Source: https://ark.intel.com. As of October 13, 2026, no other publicly available server processor with 96 or more cores offers a 5.3 GHz max frequency (Intel Xeon 6972P 96C: 3.9 GHz; NVIDIA Vera 88C: up to 3.2 GHz; Arm AGI 136C: up to 3.7 GHz; Intel Xeon 6776P with PCT: 4.6 GHz at 64C). Max boost for AMD EPYC processors is the maximum frequency achievable by any single core on the processor under normal operating conditions for server systems. Intel and Xeon are trademarks of Intel Corporation or its subsidiaries.

[7] 9xx6-063A: Comparison based on AMD internal testing and estimates and published AMD and NVIDIA specifications as of 11/05/2026. AMD EPYC scores are estimates based on AMD internal testing. System power per CPU = CPU power / 85% voltage regulator efficiency + memory power at 65% load + NIC + storage + BMC + other. 2P 6th Gen AMD EPYC 256C at 400W: 32x64GB DDR5-8000 RDIMM, GCC 16.1, estimated SPECrate®2026_int_base 1810, 1,417 W system power per 2P server (708 W per CPU). NVIDIA Vera: 88C, 450W TDP, ETL256 rack (8-CPU system), LPDDR5X-9600, GCC 15.2, estimated SPECrate®2026_int_base 925 (2P), 1,417 W system power per 2 CPUs (709 W per CPU). Within 1 MW of system power: 706 AMD EPYC 2P servers × 1,024 threads = 722,768 threads vs 706 Vera 2-CPU equivalents × 352 threads = 248,403 threads, 2.9x. Performance within 1 MW: 706 × 1810 = 1,277,550 vs 706 × 925 = 652,764 SPECrate®2026_int_base, 2.0x. Cores per rack: AMD Helion rack 32,768 cores, https://www.amd.com/en/products/rackscale-solutions/helion.html#configurations, vs NVIDIA Vera ETL256 rack 22,528 cores (256 CPUs × 88 cores), https://www.nvidia.com/en-us/data-center/products/vera-rack/. 32,768 / 22,528 = 1.45x. Agents per rack: AMD Helion rack 65,536 threads vs NVIDIA Vera ETL256 rack 45,056 threads, 2 threads per core (SMT) on both, ~1.45x. Agent counts are estimates derived from available CPU thread resources used as a proxy under a consistent theoretical workload. Actual agent capacity and throughput will vary based on workload, model, memory, software, orchestration, and system configuration. Nvidia Vera performance from: https://nvdam.widen.net/s/nmw5vblpqd/nvidia_vera_cpu_architecture_whitepaper?nvid=nv-tblg-543584. SPEC®, SPEC CPU®, and SPECrate® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. Results may vary.

[8] 9xx6-112: Comparison based on AMD internal estimates and published AMD and NVIDIA specifications as of 11/05/2026. Rack power values are AMD internal estimates. Agents per rack are the CPU threads per rack, 2 threads per core on both racks. Agents per rack kW are the agents per rack for each kW of rack power. AMD Helion rack, 48 OU liquid-cooled configuration: 144 6th Gen AMD EPYC 256-core processors at 400W (36,864 cores), 73,728 agents, 128 kW estimated rack power, 576 agents per rack kW. NVIDIA Vera ETL256 rack, 44 OU liquid-cooled configuration: 256 NVIDIA Vera 88-core CPUs at 450W TDP (22,528 cores), 45,056 agents, 185 kW estimated rack power, 244 agents per rack kW. NVIDIA specifications from https://www.nvidia.com/en-us/data-center/products/vera-rack/ and https://developer.nvidia.com/blog/inside-nvidia-vera-rubin-pod-seven-chips-five-rack-scale-systems-one-ai-supercomputer/. Results below are in the format of: [metric]. [AMD Helion rack relative to NVIDIA Vera ETL256 rack] Agents per rack, ~1.6x Agents per rack kW, ~2.4x Agent counts are estimates derived from available CPU thread resources used as a proxy under a consistent theoretical workload. Actual agent capacity and throughput will vary based on workload, model, memory, software, orchestration, and system configuration. Results may vary due to factors including system configurations, software versions and BIOS settings.

[9] 9xx6-011C: SPECrate®2026_int_base comparison based on published scores from www.spec.org as of 11/05/2026. Highest performing server CPU based on the highest published SPECrate®2026_int_base score, overall and per CPU. Generational comparison based on the highest published 2P AMD EPYC 9996 score and the highest published 2P AMD EPYC 9965 scores compiled with AOCC 5.1.0 (2.0x the performance, ~1.7x the performance / CPU W and ~1.6x the performance / CPU $) and with AOCC 5.2.0 (1.8x, ~1.5x and ~1.4x). Intel Xeon 6980P comparison based on the highest published 2P Intel Xeon 6980P score (2.3x the performance). Results and configurations below are in the format of: [processor], [cores], [Default CPU Power or TDP], [1Ku price in USD], [compiler], [score], [score / CPU W], [score / 1Ku price in USD], [Link to score] 2P AMD EPYC 9996, 512C, 500W, $14904 USD, AOCC 6.0, 2560, 4.267, 0.172, https://www.spec.org/cpu2026/results/res2026q4/cpu2026-20261014-00694.html 2P AMD EPYC 9965, 384C, 500W, $11988 USD, AOCC 5.1.0, 1270, 2.540, 0.106, https://spec.org/cpu2026/results/res2026q2/cpu2026-20260210-00037.html 2P AMD EPYC 9965, 384C, 500W, $11988 USD, AOCC 5.2.0, 1430, 2.860, 0.119, https://spec.org/cpu2026/results/res2026q2/cpu2026-20260601-00269.html 2P Intel Xeon 6980P, 256C, 500W, $13950 USD, Intel oneAPI 2026.0, 1120, 2.240, 0.080, https://spec.org/cpu2026/results/res2026q2/cpu2026-20260519-00255.html Next highest published SPECrate®2026_int_base score: 8P Intel Xeon Platinum 8490H, 1660, https://www.spec.org/cpu2026/results/res2026q3/cpu2026-20260629-00321.html. The 2P AMD EPYC 9996 score is also higher than the NVIDIA published estimate for 2P NVIDIA Vera (925), https://nvdam.widen.net/s/nmw5vblpqd/nvidia_vera_cpu_architecture_whitepaper?nvid=nv-tblg-543584 Starting with the 6th Gen AMD EPYC™ server processor family, AMD uses Default CPU Power to describe processor power consumption, succeeding AMD's historical TDP reference. Default CPU Power reflects total power consumed across the processor's compute and I/O dies for the stated performance target. Default CPU Power and TDP may both serve as processor power references for product comparison, platform planning, and performance-per-watt analysis. SPEC®, SPEC CPU® and SPECrate® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. AMD CPU prices as of 11/05/2026. Intel CPU W and prices at https://ark.intel.com/ as of 11/05/2026. Results may vary due to factors including system configurations, software versions and BIOS settings.

[10] 9xx6-031A: GROMACS performance comparison based on AMD internal testing as of 10/16/2026, results in ns/day, higher is better, mean of 3 runs (benchPEP). Each AMD EPYC system is compared with the Intel Xeon 6980P using the same memory type (DDR5-6400 RDIMM or MRDIMM-8800); 6th Gen AMD EPYC systems are also compared with the 5th Gen AMD EPYC 9755 (DDR5-6400 RDIMM). Relative to 2P Intel Xeon 6980P using the same memory type (1.00x): 9686F RDIMM-8000 1.60x, 9996 RDIMM-8000 2.80x, 9755 RDIMM-6400 1.53x, 9965 RDIMM-6400 1.82x, 9686F MRDIMM-12800 1.68x, 9996 MRDIMM-12800 2.91x. Relative to 2P 5th Gen AMD EPYC 9755 with DDR5-6400 RDIMM (1.00x): 9686F RDIMM-8000 1.04x, 9996 RDIMM-8000 1.83x, 9686F MRDIMM-12800 1.10x, 9996 MRDIMM-12800 1.96x. Results may vary due to factors including system configurations, software versions and BIOS settings. 9xx6-032A: NAMD performance comparison based on AMD internal testing as of 10/16/2026, results in ns/day, higher is better, mean of 3 runs (20stmv). Each AMD EPYC system is compared with the Intel Xeon 6980P using the same memory type (DDR5-6400 RDIMM or MRDIMM-8800); 6th Gen AMD EPYC systems are also compared with the 5th Gen AMD EPYC 9755 (DDR5-6400 RDIMM). Relative to 2P Intel Xeon 6980P with the same memory type (1.00x): 9686F RDIMM-8000 1.87x, 9996 RDIMM-8000 3.13x, 9755 RDIMM-6400 1.61x, 9965 RDIMM-6400 2.10x, 9686F MRDIMM-12800 1.87x, 9996 MRDIMM-12800 2.94x. Relative to 2P 5th Gen AMD EPYC 9755 with DDR5-6400 RDIMM (1.00x): 9686F RDIMM-8000 1.16x, 9996 RDIMM-8000 1.93x, 9686F MRDIMM-12800 1.17x, 9996 MRDIMM-12800 2.00x. Results may vary due to factors including system configurations, software versions and BIOS settings. 9xx6-034A: Quantum ESPRESSO performance comparison based on AMD internal testing as of 10/16/2026, results in wall clock time in seconds, lower is better, mean of 3 runs per benchmark. Relative performance is computed per benchmark and combined as a geometric mean across 2 benchmarks (Ausurf, Ta2O5). Each AMD EPYC system is compared with the Intel Xeon 6980P using the same memory type (DDR5-6400 RDIMM or MRDIMM-8800); 6th Gen AMD EPYC systems are also compared with the 5th Gen AMD EPYC 9755 (DDR5-6400 RDIMM). Relative to 2P Intel Xeon 6980P with the same memory type (1.00x), geometric mean: 9686F RDIMM-8000 1.41x, 9996 RDIMM-8000 1.72x, 9755 RDIMM-6400 1.25x, 9686F MRDIMM-12800 1.44x, 9996 MRDIMM-12800 1.75x. Relative to 2P 5th Gen AMD EPYC 9755 with DDR5-6400 RDIMM (1.00x), geometric mean: 9686F RDIMM-8000 1.12x, 9996 RDIMM-8000 1.38x, 9686F MRDIMM-12800 1.11x, 9996 MRDIMM-12800 1.46x. Results may vary due to factors including system configurations, software versions and BIOS settings. 9xx6-035B: WRF performance comparison based on AMD internal testing as of 11/05/2026, results in mean time per step in seconds, lower is better, mean of 3 runs (conus2.5km). Each AMD EPYC system is compared with the Intel Xeon 6980P using the same memory type (DDR5-6400 RDIMM or MRDIMM-8800); 6th Gen AMD EPYC systems are also compared with the 5th Gen AMD EPYC 9755 (DDR5-6400 RDIMM). Relative to 2P Intel Xeon 6980P with the same memory type (1.00x): 9686F RDIMM-8000 2.03x, 9996 RDIMM-8000 2.53x, 9755 RDIMM-6400 1.56x, 9686F MRDIMM-12800 2.14x, 9996 MRDIMM-12800 2.84x. Relative to 2P 5th Gen AMD EPYC 9755 with DDR5-6400 RDIMM (1.00x): 9686F RDIMM-8000 1.30x, 9996 RDIMM-8000 1.62x, 9686F MRDIMM-12800 1.51x, 9996 MRDIMM-12800 2.02x. Estimated forecast time: the measured mean time per step was extrapolated to a 3-day (72-hour) and 5-day (120-hour) forecast at the conus2.5km 15-second time step, excluding model initialization and output: 2P Intel Xeon 6980P MRDIMM-8800 7.3 and 12.2 hours; 2P AMD EPYC 9755 DDR5-6400 RDIMM 5.1 and 8.6 hours; 2P AMD EPYC 9686F MRDIMM-12800 3.4 and 5.7 hours; 2P AMD EPYC 9996 MRDIMM-12800 2.5 and 4.2 hours. Results may vary due to factors including system configurations, software versions and BIOS settings.

[12] 9xx6-047B: SPECrate®2017_int_base per-core comparison based on published scores from www.spec.org as of 11/05/2026 and AMD internal estimates as of 10/27/2026. Results below are in the format of: [processor], [cores], [score], [score / core], [Link to score] 2P AMD EPYC 9996, 512C, 5470, 10.68, https://www.spec.org/cpu2017/results/res2026q4/cpu2017-20261020-52013.html 2P AMD EPYC 9686F, 192C, AMD estimate 2890, 15.05 2P AMD EPYC 9756, 256C, AMD estimate 3300, 12.89 2P Intel Xeon 6980P, 256C, 2510, 9.80, https://www.spec.org/cpu2017/results/res2025q4/cpu2017-20251125-50522.html 1P NVIDIA Vera, 88C, AMD estimate 729.5, 8.29 1P Arm AGI, 136C, AMD estimate 972, 7.15 NVIDIA Vera and Arm AGI estimate based on Phoronix testing, https://www.phoronix.com/review/nvidia-vera-benchmarks/. AMD internal estimates and measurements are based on AMD engineering projections and subject to change. Starting with the 6th Gen AMD EPYC™ server processor family, AMD uses Default CPU Power to describe processor power consumption, succeeding AMD's historical TDP reference. Default CPU Power reflects total power consumed across the processor's compute and I/O dies for the stated performance target. Default CPU Power and TDP may both serve as processor power references for product comparison, platform planning, and performance-per-watt analysis. SPEC®, SPEC CPU® and SPECint® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. Results may vary due to factors including system configurations, software versions and BIOS settings.

[13] 9xx6-004C: SPECrate®2017_int_base comparison based on published scores from www.spec.org as of 11/05/2026 and AMD internal estimates for 2P NVIDIA Vera and 2P Arm AGI as of 10/10/2026. Highest performing 2P server CPU based on the highest published SPECrate®2017_int_base score and the highest published score per CPU (per socket). Results and configurations below are in the format of: [processor], [cores], [Default CPU Power or TDP], [1Ku price in USD], [score], [score / CPU W], [score / 1Ku price in USD], [Link to score] 2P AMD EPYC 9996, 512C, 500W, $14904 USD, 5470, 9.117, 0.367, 10.68, https://www.spec.org/cpu2017/results/res2026q4/cpu2017-20261020-52013.html 2P AMD EPYC 9965, 384C, 500W, $11988 USD, 3260, 6.520, 0.272, 8.49, https://www.spec.org/cpu2017/results/res2026q4/cpu2017-20261005-51950.html 2P Intel Xeon 6990E+, 576C, 450W, $14995 USD, 3050, 6.778, 0.203, 5.30, https://www.spec.org/cpu2017/results/res2026q2/cpu2017-20260518-51499.html 2P Intel Xeon 6980P, 256C, 500W, $13950 USD, 2510, 5.020, 0.180, 9.80, https://www.spec.org/cpu2017/results/res2025q4/cpu2017-20251126-50522.html 2P NVIDIA Vera, 176C, 450W, pricing not published, AMD estimate 1459, 3.242, n/a, 8.29 2P Arm AGI, 272C, 300W, pricing not published, AMD estimate 1944, 6.480, n/a, 7.15 NVIDIA Vera and Arm AGI scores are AMD internal estimates based on AMD engineering projections and subject to change; NVIDIA Vera estimate based on published Phoronix testing, https://www.phoronix.com/review/nvidia-vera-benchmarks/. Pricing not published for NVIDIA Vera or Arm AGI. Intel CPU W and prices at https://ark.intel.com/ as of 11/05/2026. Starting with the 6th Gen AMD EPYC™ server processor family, AMD uses Default CPU Power to describe processor power consumption, succeeding AMD's historical TDP reference. Default CPU Power reflects total power consumed across the processor's compute and I/O dies for the stated performance target. Default CPU Power and TDP may both serve as processor power references for product comparison, platform planning, and performance-per-watt analysis. SPEC®, SPEC CPU®, SPECrate® and SPECint® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. AMD CPU prices as of 11/05/2026. Results may vary due to factors including system configurations, software versions and BIOS settings.

[14] 9xx6-062A: Based on AMD internal testing as of 11/05/2026, geomean IPC uplift at the same fixed frequency (3 GHz) of an AMD EPYC™ 9006 Series 9996 CPU compared to an AMD EPYC™ 9005 Series 9B45 CPU, each with 16 cores and 16 threads active, using a select set of 27 workloads (75 tests) in 7 categories: server kernels, AI, database, HPC, cloud native, web serving, and industry-standard benchmarks. Industry-standard benchmarks include estimated SPECrate®2017_int_base, estimated SPECrate®2017_fp_base, estimated SPECjbb®2015-Composite max-jOPS, and an enterprise Java workload based on SPECpower_ssj®2008. Tests of the same workload are combined by geometric mean into a workload result. The overall results are combined by weighted geometric mean into a category result. The overall result is the geometric mean of the 7 category results. Methods of measuring IPC vary generation to generation based on the most current and relevant workloads, so IPC uplifts are only comparable between adjacent AMD EPYC server processor generations. AMD EPYC 9996: 1P, 16 cores active in 1 CCD, fixed 3 GHz, DDR5-8000 2Rx4 64GB DIMMs, BIOS TNIV0090D. AMD EPYC 9B45: 1P, 16 cores active in 2 CCDs (8 per CCD), fixed 3 GHz, DDR5-6400 2Rx4 64GB DIMMs, BIOS TVDT1005B. Both: Ubuntu 24.04.2 LTS, kernel 6.8.0-38-generic, performance CPU governor. Results (9996 relative to 9B45, 1.00x = parity): Overall 1.27x (27% uplift) Server kernels (5 workloads): 1.31x AI (5 workloads): 1.28x HPC (8 workloads): 1.27x Cloud native (1 workload): 1.25x Web serving (1 workload): 1.22x Industry-standard benchmarks (4 workloads): 1.20x Results may vary due to factors including system configurations, software versions and BIOS settings. SPEC®, SPEC CPU®, SPECrate®, SPECjbb®, and SPECpower_ssj® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. Java is a registered trademark of Oracle and/or its affiliates.

[15] 9xx6-056A: Estimated rack-level throughput uplift within a 100 kW rack (Dell IR7000), based on AMD internal modeling across six workloads (SPECrate®2026_int_base, SPECjbb®2015 MultiJVM max-jOPS, NGINX, Redis, memcached, TPROC-C MySQL), as of 11/05/2026. SPECrate®2026_int_base and SPECjbb®2015 scores for AMD EPYC and Intel Xeon are published results from www.spec.org as of 11/05/2026; NVIDIA Vera SPECrate®2026_int_base is the NVIDIA published estimate (https://nvdam.widen.net/s/nmw5vblpqd/nvidia_vera_cpu_architecture_whitepaper?nvid=nv-tblg-543584); other NVIDIA Vera scores and the NGINX, Redis, memcached and TPROC-C results are AMD measurements and estimates. AMD internal estimates and measurements are based on AMD engineering projections and subject to change. Rack model: NVIDIA Vera 88C 450W (176 CPUs/rack, 2S), Intel Xeon 6980P 128C 500W (150 CPUs/rack, 2S), 6th Gen AMD EPYC 9996 256C 600W (126 CPUs/rack, 2S). Published results: [processor], [SPECrate®2026_int_base], [Link], [SPECjbb®2015 MultiJVM max-jOPS], [Link] 2P AMD EPYC 9996, 2560, https://www.spec.org/cpu2026/results/res2026q4/cpu2026-20261014-00694.html, 2746272, https://www.spec.org/jbb2015/results/res2026q4/jbb2015-20261013-01765.html 2P AMD EPYC 9965, 1430, https://www.spec.org/cpu2026/results/res2026q2/cpu2026-20260601-00269.html, 1385555, https://www.spec.org/jbb2015/results/res2025q1/jbb2015-20250319-01562.html 2P Intel Xeon 6980P, 1120, https://www.spec.org/cpu2026/results/res2026q2/cpu2026-20260519-00255.html, 1115670, https://www.spec.org/jbb2015/results/res2026q3/jbb2015-20260818-01736.html Results (rack throughput, normalized to NVIDIA Vera): NVIDIA Vera 88C baseline 1.00 2P Intel Xeon 6980P (128C) 1.38 2P AMD EPYC 9965 (192C) 2.30 2P AMD EPYC 9996 (256C) 3.48 TPROC-C workload is an open-source workload derived from TPC-Benchmark™ Standard, and as such is not comparable to published TPC-C™ results, as the results do not comply with the TPC-C Benchmark Standard. TPC, TPC Benchmark and TPC-C are trademarks of the Transaction Processing Performance Council. SPEC®, SPECrate® and SPECjbb® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. Results may vary due to factors including system configurations, software versions and BIOS settings.

[16] 9xx6-021: Rack-level general purpose performance comparison within a 100 kW rack power constraint based on AMD analysis as of July 2026 comparing 2P AMD EPYC™ "Venice" 9996 (512 total cores) vs 2P Nvidia Vera (176 total cores, Olympus). Performance reflects geometric mean across 6 workloads: Server-Side Java® Multi-JVM Max jOPS, Web Serving (NGINX/WRK), Key-Value Store (Redis), In-memory caching (Memcached), and Transactional DB (TPROC-C). All platforms use 2-socket nodes. System manufacturers may vary configurations, yielding different results. See https://www.amd.com/content/dam/amd/en/documents/solutions/ai/methodology-description.pdf for full derivations and assumptions. TPROC-C workload is an open-source workload derived from TPC-Benchmark™ Standard, and as such is not comparable to published TPC-C™ results, as the results do not comply with the TPC-C Benchmark Standard. TPC, TPC Benchmark and TPC-C are trademarks of the Transaction Processing Performance Council. Preliminary performance estimates based on AMD engineering projections or measurements as of July 22, 2026 and subject to change. SPEC® and SPECint® are registered trademarks of the Standard Performance Evaluation Corporation. Learn more at spec.org.

© 2026 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. Other product names used in this publication are for identification purposes only and may be trademarks of their respective owners.

AMD
together we advance_