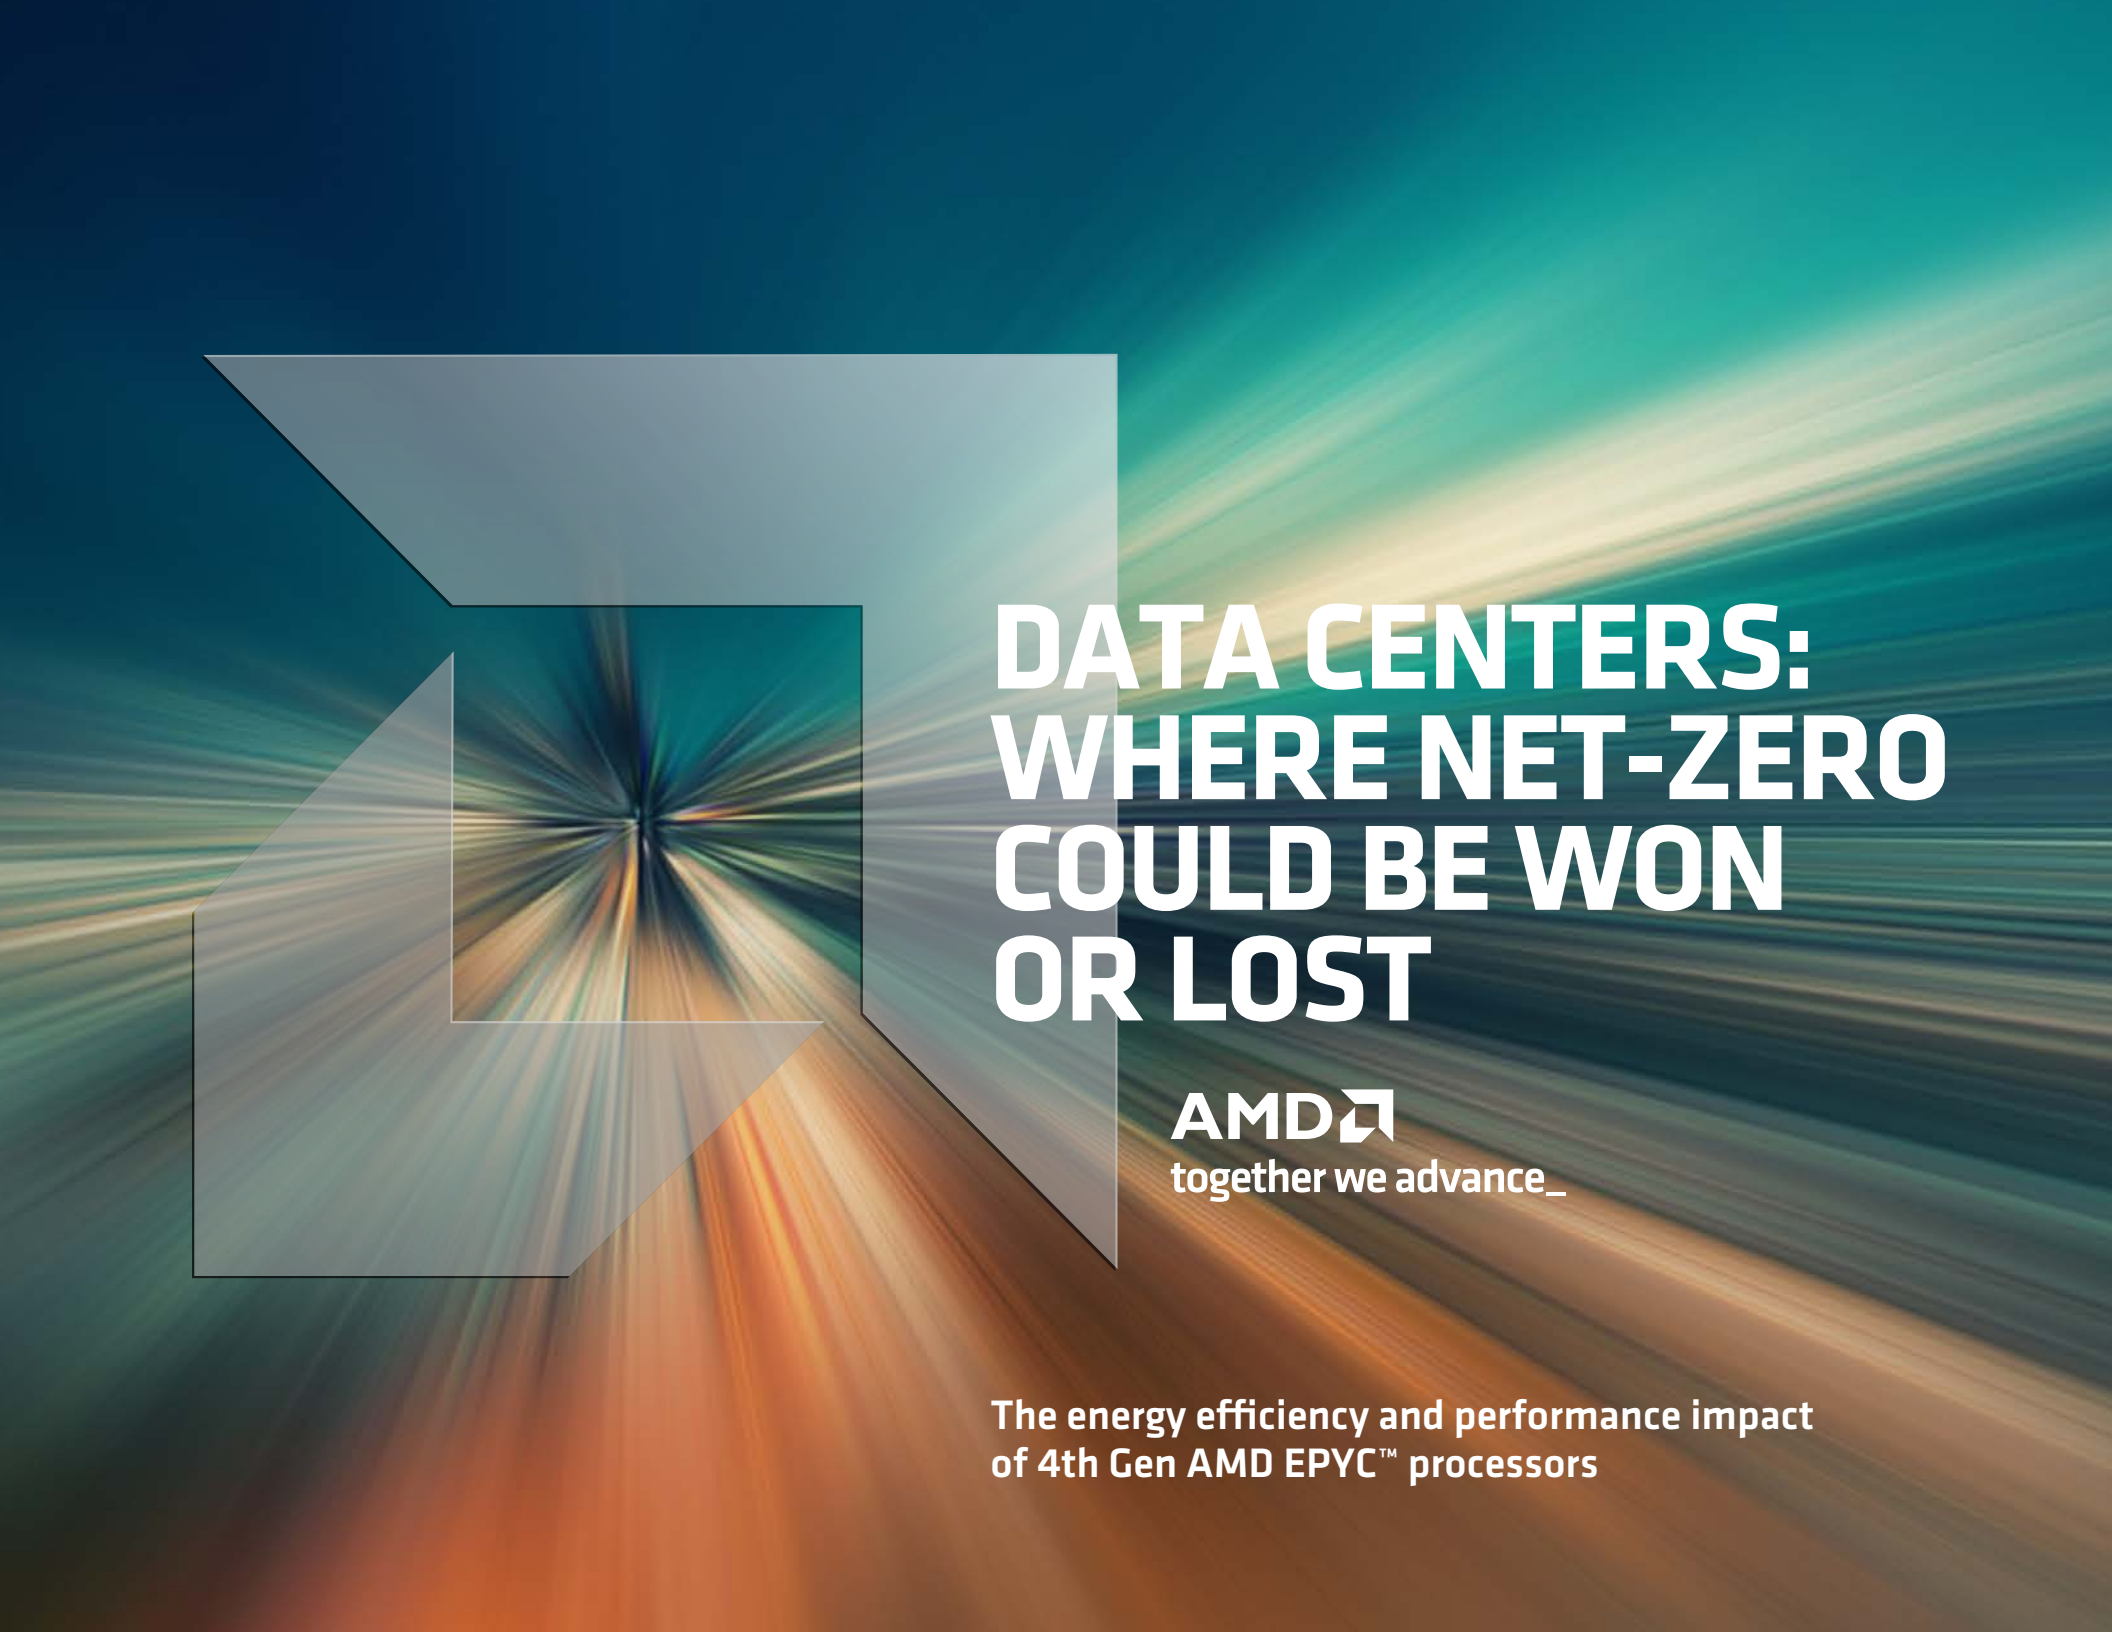

DATA CENTERS: WHERE NET-ZERO COULD BE WON OR LOST

AMD
together we advance_

The energy efficiency and performance impact of 4th Gen AMD EPYC™ processors

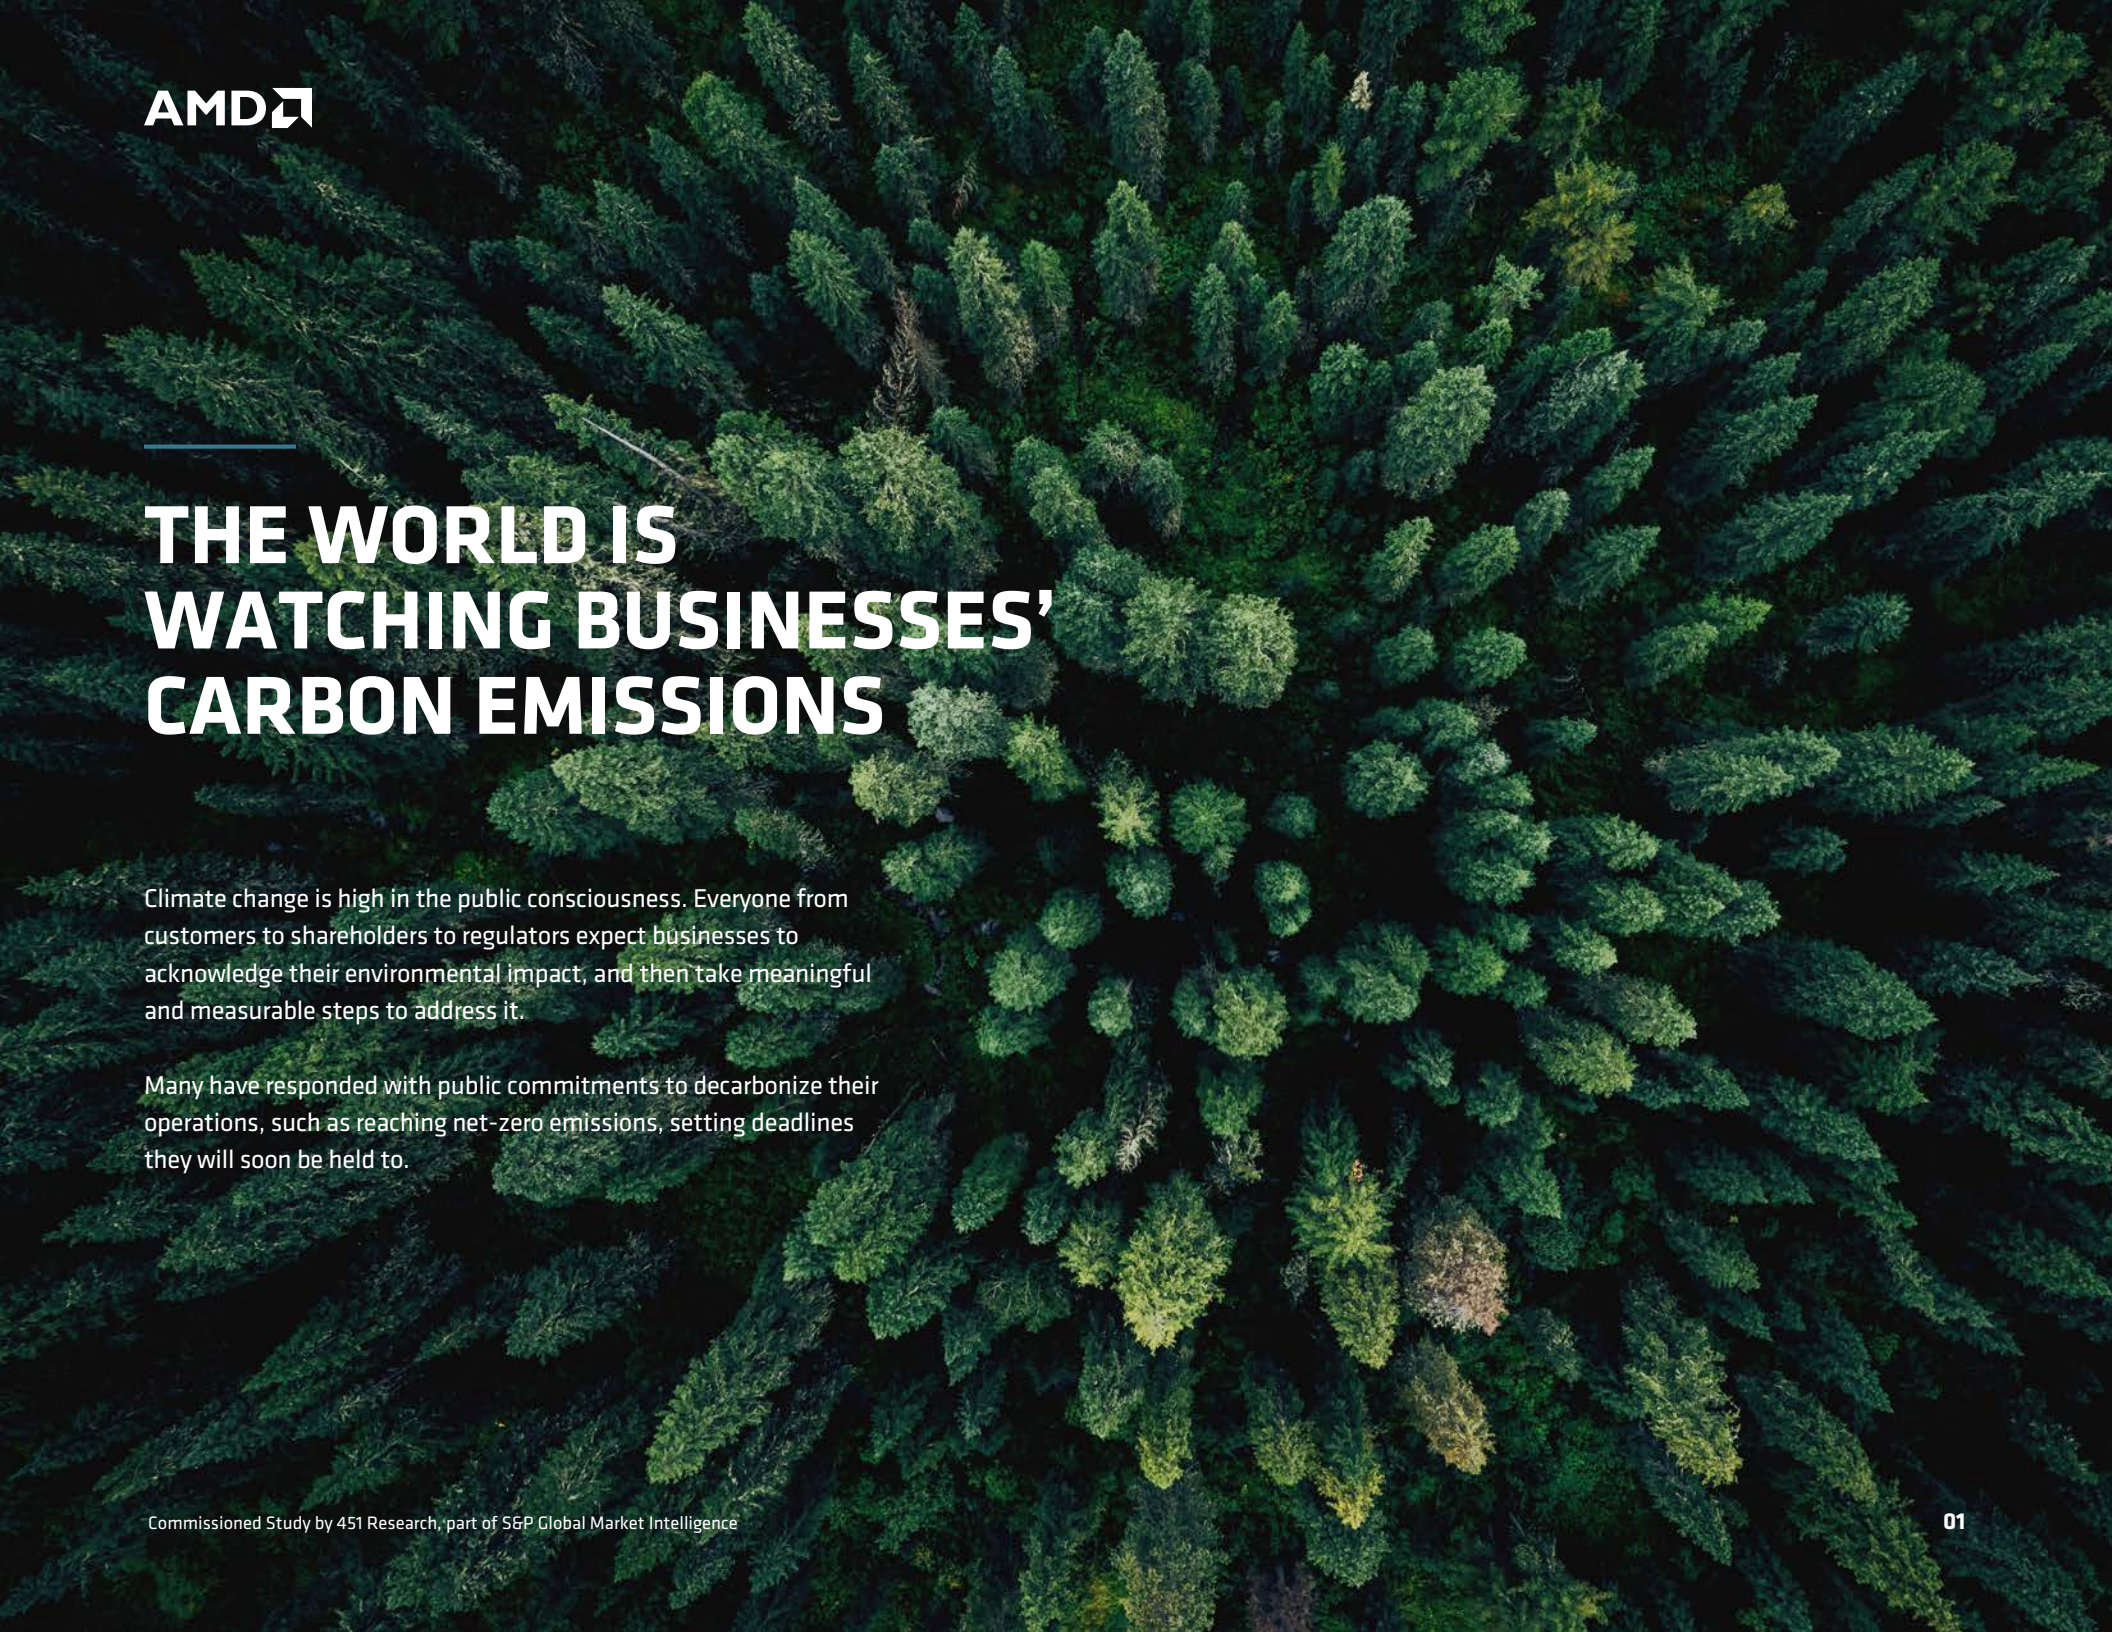

THE WORLD IS WATCHING BUSINESSES' CARBON EMISSIONS

Climate change is high in the public consciousness. Everyone from customers to shareholders to regulators expect businesses to acknowledge their environmental impact, and then take meaningful and measurable steps to address it.

Many have responded with public commitments to decarbonize their operations, such as reaching net-zero emissions, setting deadlines they will soon be held to.

Commissioned Study by 451 Research, part of S&P Global Market Intelligence

KEY NUMBERS:

SERVER AND DATA CENTER ENERGY CONSUMPTION

DATA CENTERS CONSUME...

2% OF **1.1-1.5%**

2% of electricity in the US[1]

1.1-1.5% all electricity globally[2]

IN EUROPE, DATA CENTER ENERGY USE IN 2025 IS FORECASTED TO BE:

21%

above 2018 levels[3]

~1/2

Nearly half of IT decision-makers report that IT operations now account for most (25%) or all (19%) of their environmental impact

2.2%

Compound annual growth rate of data center energy consumption in the US since 2019, compared to 0.3% total growth in demand for power[4]

AMD OFFERS LEADERSHIP ENERGY EFFICIENCY

It's little wonder that IT leaders are adding energy efficiency to their data center modernization criteria, alongside the usual requirements of high performance, robust security, and ample flexibility.

4th Gen AMD EPYC™ processors are the latest in the AMD line of server processors designed for performance and efficiency, to help reduce enterprise energy usage, save costs, and help companies meet their data center sustainability goals.

With a core infrastructure based on AMD processors, businesses can take real action with measurable results, delivering targeted workloads while helping reduce energy consumption and cost. All without compromising performance.

4TH GEN AMD EPYC™ PROCESSOR VERSUS THE COMPETITION

IT'S GOOD FOR BUSINESS AND HELPS ENTERPRISES MEET SUSTAINABILITY GOALS[5,6]

SERVERS REQUIRED TO RUN 2,000 VIRTUAL MACHINES (1 CORE, 8GB MEMORY):

17 Intel
2P Intel Platinum 8490H

11 AMD
2P AMD EPYC™ 9654

RESULTING IN 35% FEWER SERVERS NEEDED, REQUIRING UP TO AN ESTIMATED 29% LESS POWER ANNUALLY.

ESTIMATED **46%**
Lower CAPEX

BENEFITS TO ORGANIZATIONS CHOOSING 4TH GEN AMD EPYC™ PROCESSORS:

- ENERGY EFFICIENCY
- DATA CENTER CONSOLIDATION
- LOW TCO

IMAGINE SWAPPING OUT

1M FOR **318,000**

Intel 8380 servers

AMD 9654 servers

TO DELIVER 133 MILLION VMS AND... PREVENT ~10 MILLION METRIC TONS OF RELATED CO2 EMISSIONS FROM ENTERING THE ATMOSPHERE OVER 3 YEARS

THAT'S EQUIVALENT TO THE AMOUNT OF CO2 REMOVED FROM THE ATMOSPHERE BY...

4M

acres of U.S. forest annually
(An area of forest larger than Yellowstone National Park.)

THE AMD COMMITMENT: 30X25

Commercial energy prices continue to soar worldwide. Businesses must become more efficient to save energy and power costs.

The choice of processor for enterprise servers can make a difference with energy efficiency, whether for new deployments or refreshing servers already in the data center.

AMD has a longstanding commitment to improving energy efficiency, across its product lines, to enable businesses to reduce power usage, greenhouse gas emissions, and costs. But doing so without compromising on computing capability and application workload throughput.

AMD has a goal to deliver a 30x increase in energy efficiency by 2025, using a 2020 baseline, for accelerated data center compute nodes running high-performance computing and AI workloads.[7] Accomplishing this ambitious goal equates to a 97 percent reduction in energy use per computation from 2020-2025, and requires AMD to increase the energy efficiency of a compute node at a rate that is more than 2.5x faster than the aggregate industry-wide improvement made during the last five years.[8]

THE EFFICIENCY OF 4TH GEN EPYC™ PROCESSORS IS JUST THE NEXT CHAPTER OF THIS COMMITMENT.

[1] https://www.energy.gov/eere/buildings/data-centers-and-servers

[2] https://energyinnovation.org/2020/03/17/how-much-energy-do-data-centers-really-use/. © Energy Innovation®, used under Creative Commons License 4.0 (CC BY 4.0)

[3] European Commission, Directorate-General for Communications Networks, Content and Technology, Montevecchi, F., Stickler, T., Hintemann, R., et al., Energy-efficient cloud computing technologies and policies for an eco-friendly cloud market : final study report, Publications Office, 2020, https://data.europa.eu/doi/10.2759/3320, page 16

[4] https://www.amd.com/en/campaigns/lowering-the-impact

[5] SP5TCO-036: As of 01/18/2023 based on AMD Internal analysis using the AMD EPYC™ Server Virtualization & Greenhouse Gas Emission TCO Estimation Tool - version 12.10 estimating the cost and quantity of 2P AMD 96 core EPYC™ 9654 powered server versus 2P Intel® Xeon® 60 core Platinum 8490H based server solutions required to deliver 2000 total virtual machines (VM), requiring 1 core and 8GB of memory per VM for a 3-year period. This includes VMware software license cost of $6,558.32 per socket + one additional software for every 32 CPU core increment in that socket.Environmental impact estimates made leveraging this data, using the Country / Region specific electricity factors from the '2020 Grid Electricity Emissions Factors v1.4 – September 2020', and the United States Environmental Protection Agency 'Greenhouse Gas Equivalencies Calculator'.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. For additional details, see https://www.amd.com/en/claims/epyc4#SP5TCO-036.

[6] SP5TCO-021A: This analysis is a theoretical exercise, made as of 11/15/2022 based on AMD Internal analysis using the AMD EPYC™ Bare Metal Server & Greenhouse Gas Emission TCO Estimation Tool - version 6.40 estimating the cost and quantity of 2P AMD EPYC™ 9654 (96 core/CPU) powered server versus 2P Intel® Xeon® Gold 8380 (40 core/CPU) based server solutions required to deliver 133 million total virtual machines (VM) based on VMmark tiles in published results. Environmental impact estimates made leveraging this data, using the Country / Region specific electricity factors from the '2020 Grid Electricity Emissions Factors v1.4 – September 2020', and the United States Environmental Protection Agency 'Greenhouse Gas Equivalencies Calculator'. 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. For additional details, see https://www.amd.com/en/claims/epyc4#SP5TCO-021A.

[7] Includes high-performance CPU and GPU accelerators used for AI training and High-Performance Computing in a 4-Accelerator, CPU hosted configuration. Goal calculations are based on performance scores as measured by standard performance metrics (HPC: Linpack DGEMM kernel FLOPS with 4k matrix size. AI training: lower precision training-focused floating-point math GEMM kernels such as FP16 or BF16 FLOPS operating on 4k matrices) divided by the rated power consumption of a representative accelerated compute node including the CPU host + memory, and 4 GPU accelerators.

[8] Includes high-performance CPU and GPU accelerators used for AI training and High-Performance Computing in a 4-Accelerator, CPU hosted configuration. Goal calculations are based on performance scores as measured by standard performance metrics (HPC: Linpack DGEMM kernel FLOPS with 4k matrix size. AI training: lower precision training-focused floating-point math GEMM kernels such as FP16 or BF16 FLOPS operating on 4k matrices) divided by the rated power consumption of a representative accelerated compute node including the CPU host + memory, and 4 GPU accelerators.

C2946521-0225