Boilerplate language for OEMs to add to their AMD bids

AMD has a proven track-record of prioritizing and delivering on environmental sustainability initiatives and performance, backed by over 25 years of public reporting. As designers of microprocessors, AMD embraces the responsibility to help protect the planet and the opportunity to help others save energy and reduce greenhouse gas emissions.

AMD environmental programs and initiatives extend across its value chain, including operations, supply chain manufacturing and products. The company has set ambitious goals through 2025/2030 that include bold advancements in energy efficiency for accelerated computing applications; setting a science-based GHG emissions reductions goal for its operations (aligned with a 1.5 degree Celsius scenario); and working with suppliers to increase efficient use of resources and renewable energy.

AMD has received numerous sustainability recognitions including 3BL Media's 100 Best Corporate Citizens (2021), CDP's Supplier Engagement Leader (2021), and Newsweek's Most Responsible Companies (2022).

AMD Environmental Sustainability Overview

Our Approach to Environmental Sustainability

As designers of microprocessors during a period of amazing growth in technology, AMD embraces the responsibility to help protect our planet and the opportunity to help others save energy and reduce greenhouse gas (GHG) emissions.

Our environmental programs and initiatives span our operations, supply chain and products. And we set ambitious goals and publicly report on our progress annually.

Our fundamental approach to environmental sustainability is based on three pillars:

- ***MINIMIZE*** environmental impacts at AMD and in our supply chain,
- ***INNOVATE*** on collaborative solutions to address environmental challenges, and
- ***ADVANCE*** the sustainability performance for customers.

Learn More

CDP SUPPLIER ENGAGEMENT LEADER 2021

3BL MEDIA 100 BEST CORPORATE CITIZENS 2021

Corporate ESG Performance Prime RATED BY ISS ESG

FTSE4Good

RBA COMPASS AWARD 2020 INNOVATION

AMD

How AMD & Our Partners Are Advancing Environmental Sustainability

Minimize Environmental Impacts
in AMD & Our Supply Chain

Innovate on Collaborative Solutions
to Environmental Challenges

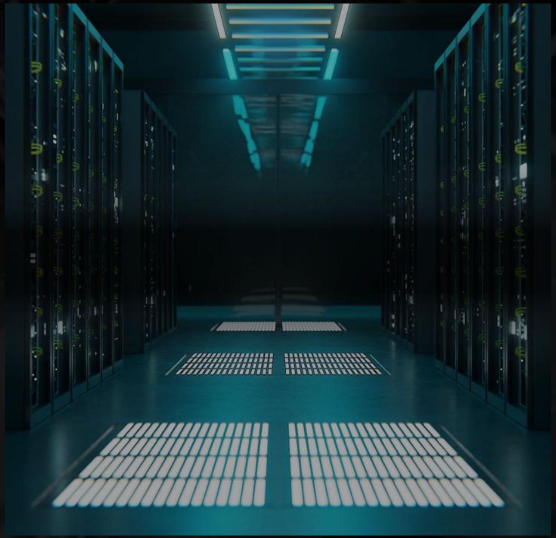

Advance Sustainability
Performance for Our Customers

Stand Alone Slide (if needed)

Established Environmental Performance & Goals

Operations

COMPLETE: **38%** reduction in greenhouse gas (GHG) emissions from AMD operations (2014-2020)

GOAL: **50%** reduction in GHG emissions from AMD operations (2020 – 2030)

Supply Chain

COMPLETE: **73%** lower direct GHG emissions for AMD wafer production compared to industry average (2020)

GOAL: **100%** of AMD manufacturing suppliers have public GHG emissions reduction goals and 80 percent source renewable energy by 2025

Product Use

COMPLETE: **31.7x** increase in energy efficiency for mobile processors (2014-2020)

GOAL: **30x** increase in energy efficiency for AMD processors and accelerators powering servers for AI-training and HPC (2020-2025)

SOURCE: AMD.COM/CORPORATERESPONSIBILITY/ENVIRONMENTAL-STEWARDSHIP

AMD Operations & Supply Chain

- Our operational strategy includes sourcing renewable energy, harvesting rainwater, composting/recycling waste and engaging employees in our "Go Green" Program.
- Looking ahead, our operational goal is a **50% absolute** reduction in GHG emissions from AMD operations (2020-2030).

- Our supply chain strategy includes engaging all AMD manufacturing suppliers on expectations and performance, as well as tracking key improvement metrics for our largest suppliers (80%+ by spend).
- Looking ahead to 2025, our supply chain goals include **100%** of AMD manufacturing suppliers having a public GHG emission reduction goal, **80%** sourcing renewable energy, and **100%** of supplier factories having a Responsible Business Alliance audit

AMD Operations (2014-2020)

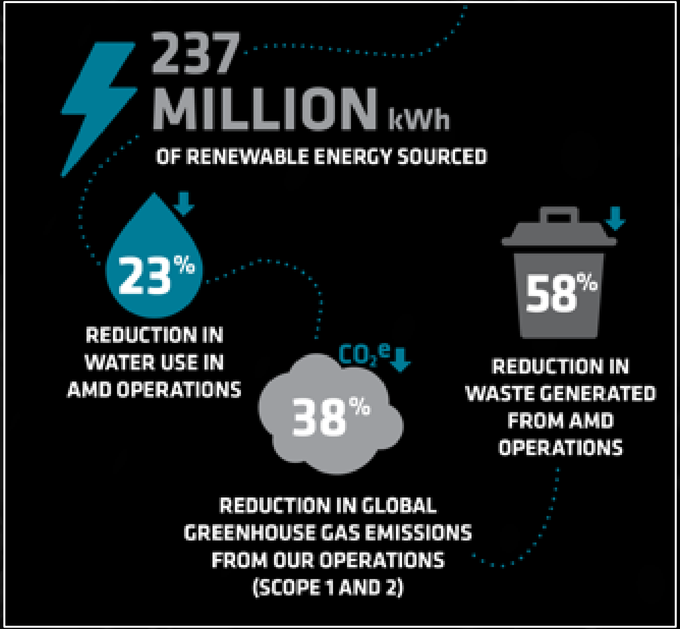

SOURCE: AMD.COM/CORPORATERESPONSIBILITY/ ENVIRONMENTAL-STEWARDSHIP

AMD

Collaborative Solutions to Environmental Challenges

Technology alone cannot achieve societal progress. Rather, it’s the people and collaborations that put high-performance computing to work and spark innovations that benefit society.

That is why we work with our customers, industry and other stakeholders to advance scientific research and public policy in ways that serve the greater good.

Scientists are advancing research on climate change and the impacts of extreme weather using AMD-powered servers.

Through strategic partnerships, AMD technology is enabling the optimization of renewable energy production.

We work with industry, government, civil society and other stakeholders to advance the UN Sustainable Development Goals (SDGs).

Needs input from Sales/Marketing; leads into next slide with customer-facing proof points

Advancing Sustainability Performance for Customers

- Understanding our customers' sustainability goals and sharing their visions are critical to how we operate. We incorporate these insights into our key product innovations to help address the challenges and opportunities ahead.
- Maximize computing performance delivered per watt of energy consumed means AMD solutions can increase enterprise IT energy efficiency and data center workload optimization while reducing energy use and greenhouse gas (GHG) emissions.

Transformed development of marine propellers to be more efficient and environmentally friendly

LINK

Tripled data center performance with lower energy consumption in the same footprint

LINK

Reduced power consumption 30-40% compared to cloud services

LINK

Reduced the number of servers and allowed for expansion in the data center

LINK

Shrank the data center footprint by 25 percent

LINK

AMD

Leveraging Innovation Across the Product Portfolio

LEADERSHIP PERFORMANCE PER WATT

ADVANCED POWER MANAGEMENT

CUTTING-EDGE PROCESS TECHNOLOGY

HIGH CORE DENSITY

EFFICIENTLY INTEGRATED MEMORY & CPU

OPTIMIZED WORKLOADS

AMD RYZEN™ PROCESSORS ADVANCING SUSTAINABILITY

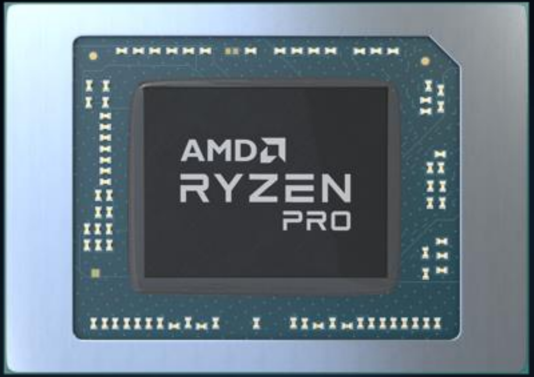

ENERGY EFFICIENT LAPTOPS WITHOUT COMPROMISE

- 50% less typical energy use from 2018-2022[1]
- 2.5X more performance vs. 4-year-old laptop[2]
- AMD exceeding Energy Star 8.0 requirements by 41% in 2022[3]

POWERFUL, EFFICIENT DESKTOP PROCESSING

- Yields up to 2x more performance to power ratio than the competition
- Up to 67% lower power per core
- Up to 39% faster rendering

See endnote CGP-24

COLLABORATION WITH ENERGY REGULATIONS

- California Energy Commission (CEC)
- US EPA ENERGY STAR
- European Union Eco-design for computers (Lot 3) and servers (Lot 9)
- China National Institute of Standardization (CNIS)

Not a comprehensive list

AMD

ESTIMATED SAVINGS OVER 4 YEARS ACROSS 10K LAPTOPS

UPDATING RYZEN 2000 SERIES TO RYZEN 5000 SERIES

272,000

Kilowatt Hours[1]

183 METRIC TONS

CO2 reduced[2]

3,187

Tree seedlings grown[2]

2.4X

Performance increase[3]

1Based on Energy Star measurements of Ryzen 2500 vs. Ryzen 5800 as measure in AMD lab. 2Estimate Based on Energy Saved by Energy Star calculations across 10K Units using EPA calculatorhttps://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator. 3Based on 50:50 weighted avg increase in Cinebench 15nt, 3DMark across Ryzen 2500 to Ryzen 5800.

AMD

Advancing Data Center Sustainability

Modern data centers are continuously striving for greater efficiency and scalability while delivering increased performance and security.

As designers of cutting-edge server CPUs and GPUs, we recognize our important role in addressing these critical priorities.

We are focused on accelerating server energy efficiency and delivering high-performance computing to help tackle some of the world’s toughest challenges, including climate change research.

Our goal is a 30x increase in energy efficiency for AMD processors and accelerators for AI-training and HPC from 2020-2025.[4]

Accelerating Data Center Sustainability

Goal: 30x increase in energy efficiency for AMD processors and accelerators for AI-training and HPC from 2020-2025

This represents more than a 2.5x acceleration of the industry trends from 2015-2020 as measured by the worldwide energy consumption for these computing segments[5] and equates to a 97% reduction in energy use per computation.

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

> *"The energy efficiency goal set by AMD for accelerated compute nodes used for AI training and High-Performance Computing fully reflects modern workloads, representative operating behaviors and accurate benchmarking methodology."*
>
> ~ Dr. Jonathan Koomey, President, Koomey Analytics

Accelerating Data Center Sustainability

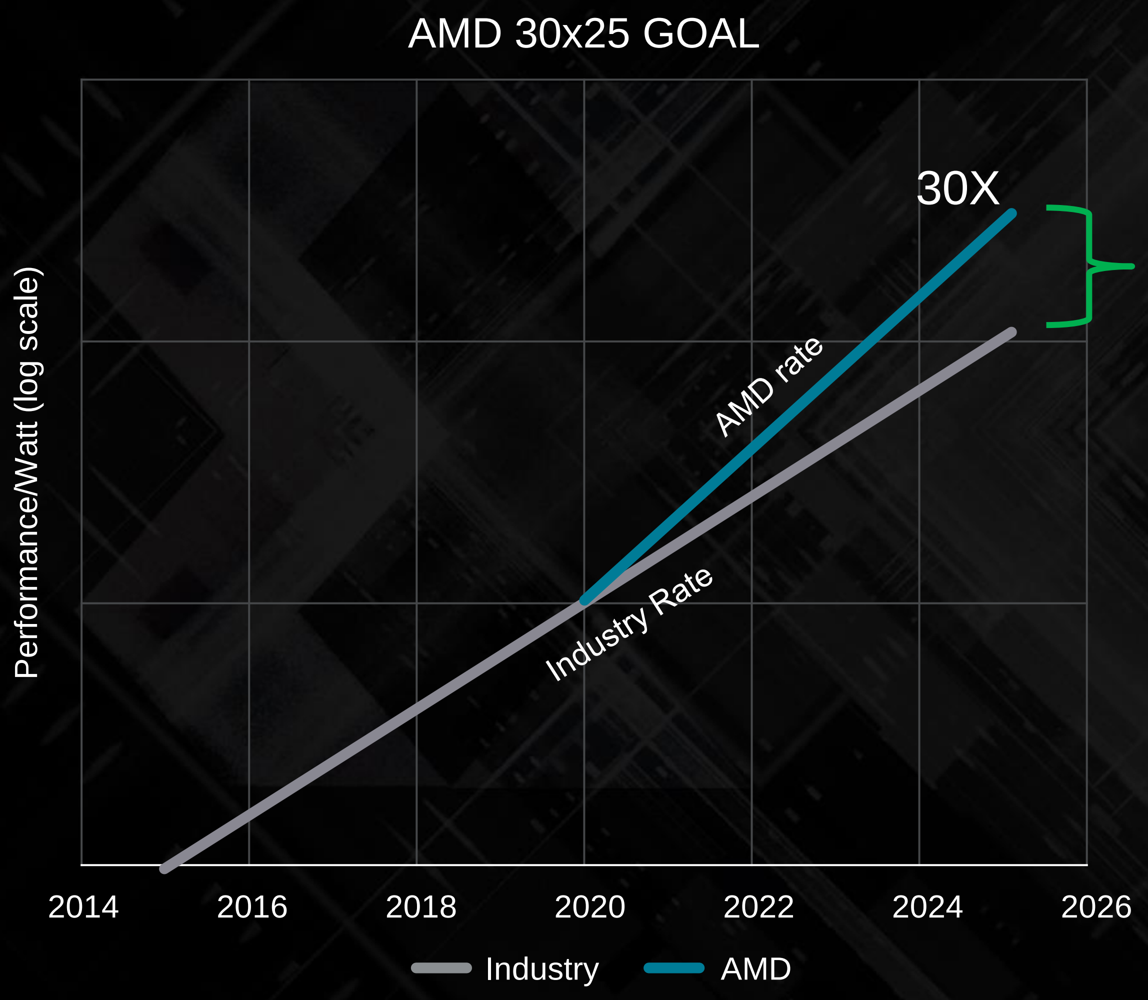

51.4 billion kilowatt-hours (kwh) of electricity could be saved from 2021-2025 relative to baseline trends if all AI and HPC server nodes globally made similar gains, equivalent to: [6]

- ❑ $6.2 billion USD in electricity costs
- ❑ 36.4 million metric tons of carbon dioxide equivalent (MTCO2e)
- ❑ Electricity generated by 7,500 wind turbines in a year, or
- ❑ Carbon sequestered by 600 million tree seedlings grown for 10 years

AMD

AMD EPYC[TM] processors power the most energy efficient x86 servers, delivering exceptional performance and reducing energy costs.

EPYC CPUs help minimize environmental impacts from data center operations while advancing your company's sustainability goals and objectives.

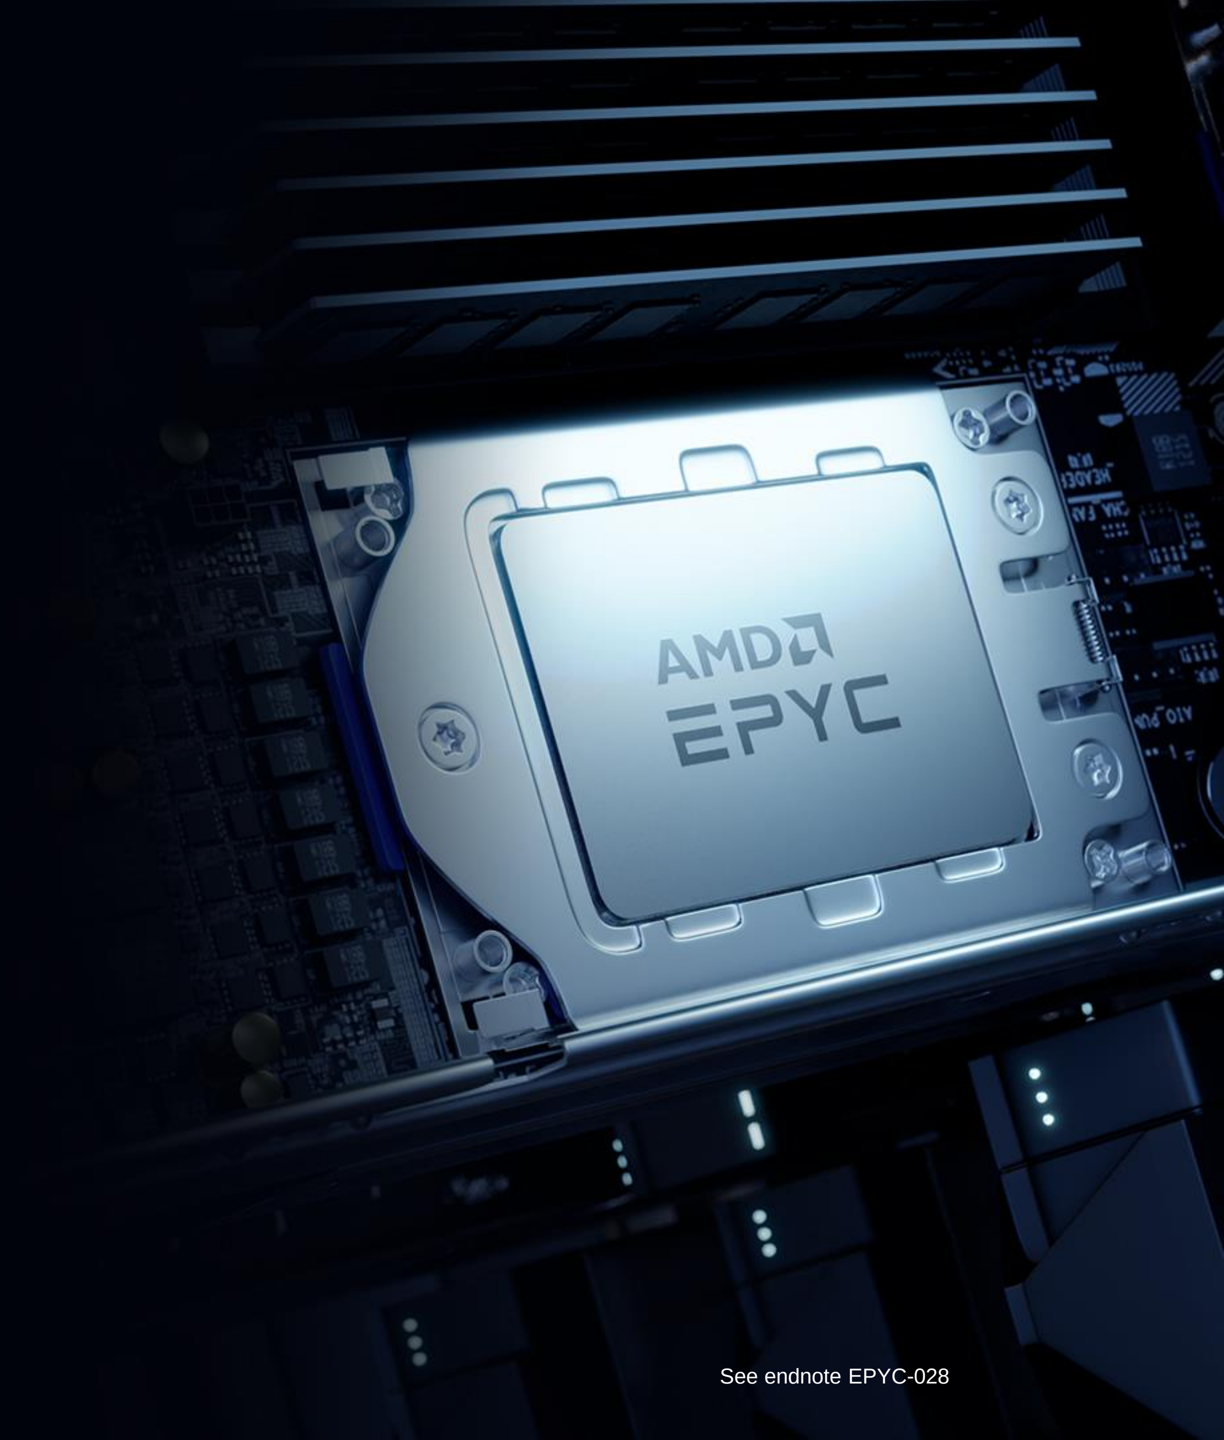

See endnote EPYC-028

Fewer Servers Use Less Power, Leading to Lower GHG Emissions

10,000 Integer Score*

Intel Scalable Platinum 8380 vs AMD EPYC™ 7763

INTEL®

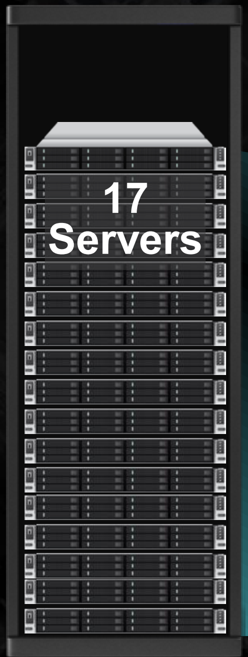

Intel Platinum 40c 8380
SPEC score 602
Link to SPEC result

EPYC Savings

29% (up to) Fewer Servers

29% (up to) Less Power

~41 Acres of Forest Annually,
equivalent carbon sequestration**

plus 30% (up to) Lower 3yr TCO*

AMD EPYC™

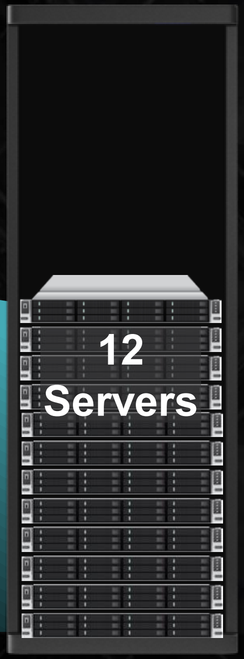

AMD EPYC 64c 7763
SPEC score 861
Link to SPEC result

AMD CPU pricing Jan 2022. Intel® Xeon® Scalable CPU pricing from https://ark.intel.com as of Jan 2022. All pricing is in USD.
All performance scores based on AMD internal testing, Oct 2021. Analysis based on AMD EPYC Bare Metal Server and Greenhouse Gas TCO Tool, v4.2.
* Comparisons made using SPECrate®2017_int_base on 01/14/2022. SPEC®, SPECrate® and SPEC CPU® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information.
** Values are for USA.

See endnote MLNTCO-020

AMD

Power of High Core Density

320 Virtual Machines (VMs)

INTEL®

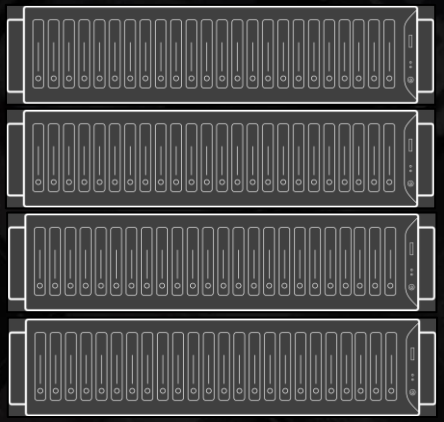

4x 2-socket
Platinum 8380 CPU
(40 core)

AMD

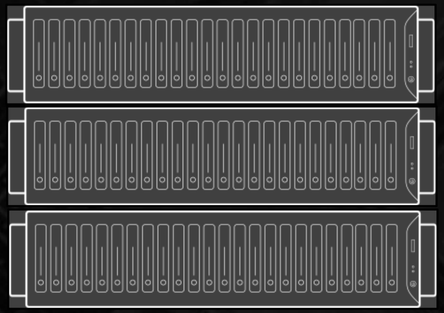

3x 2-socket
EPYC™ 7713 CPU
(64 core)

SOLUTION

320 VMs

1 core / VM, with 8GB memory / VM

VMware® vSphere Enterprise Plus w/ Production Support

AMD EPYC CPUs

THE CLEAR WINNER

23% up to — Lower Greenhouse Gas Emissions*

12% up to — Lower total 3-year TCO cost

28% up to — Lower hardware acquisition cost

* As determined by AMD GHG CO2 Emissions Calculator v2.30
See endnotes: MLNTCO-008a

AMD

AMD

AMD Instinct™ accelerators are engineered from the ground up for this new era of data center computing, supercharging HPC and AI workloads to propel new discoveries.

The AMD Instinct™ family of accelerators deliver industry leading performance from single server solutions up to the world’s largest supercomputers used to tackle our most pressing challenges.

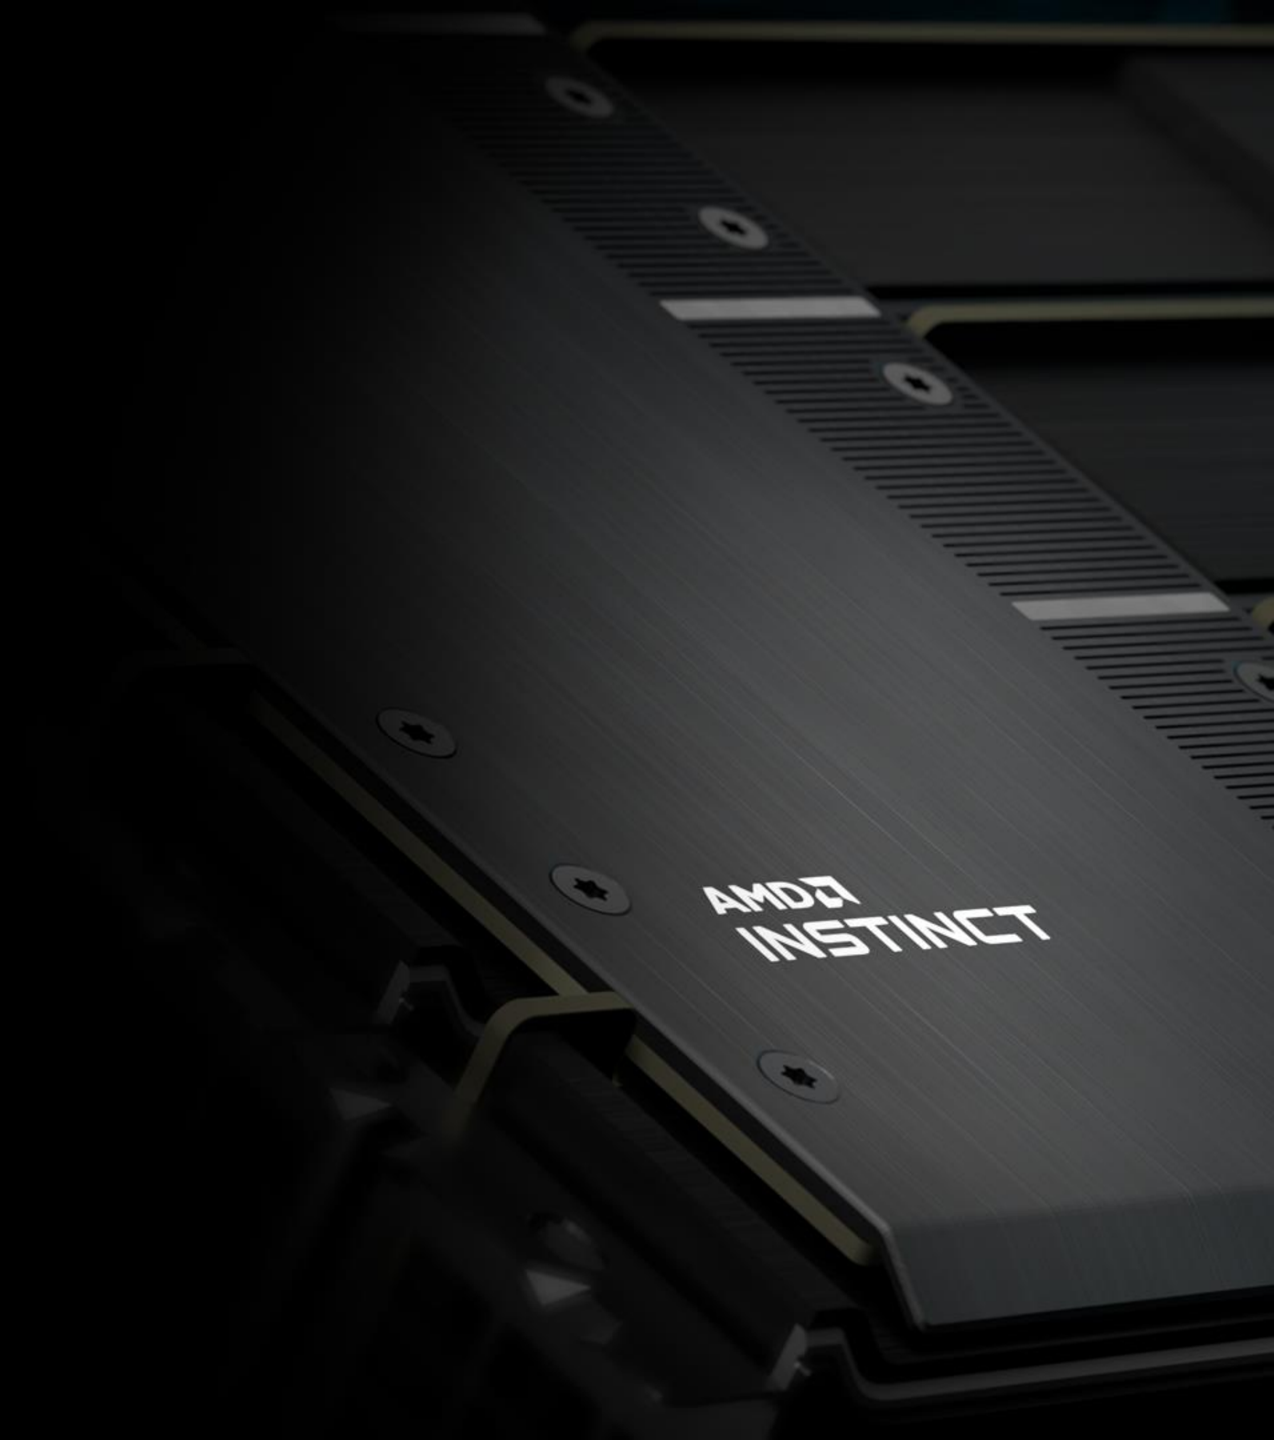

AMD Instinct MI200 Series
Up to 5.0X better Peak Perf/Watt Than Previous Generation For HPC Workloads

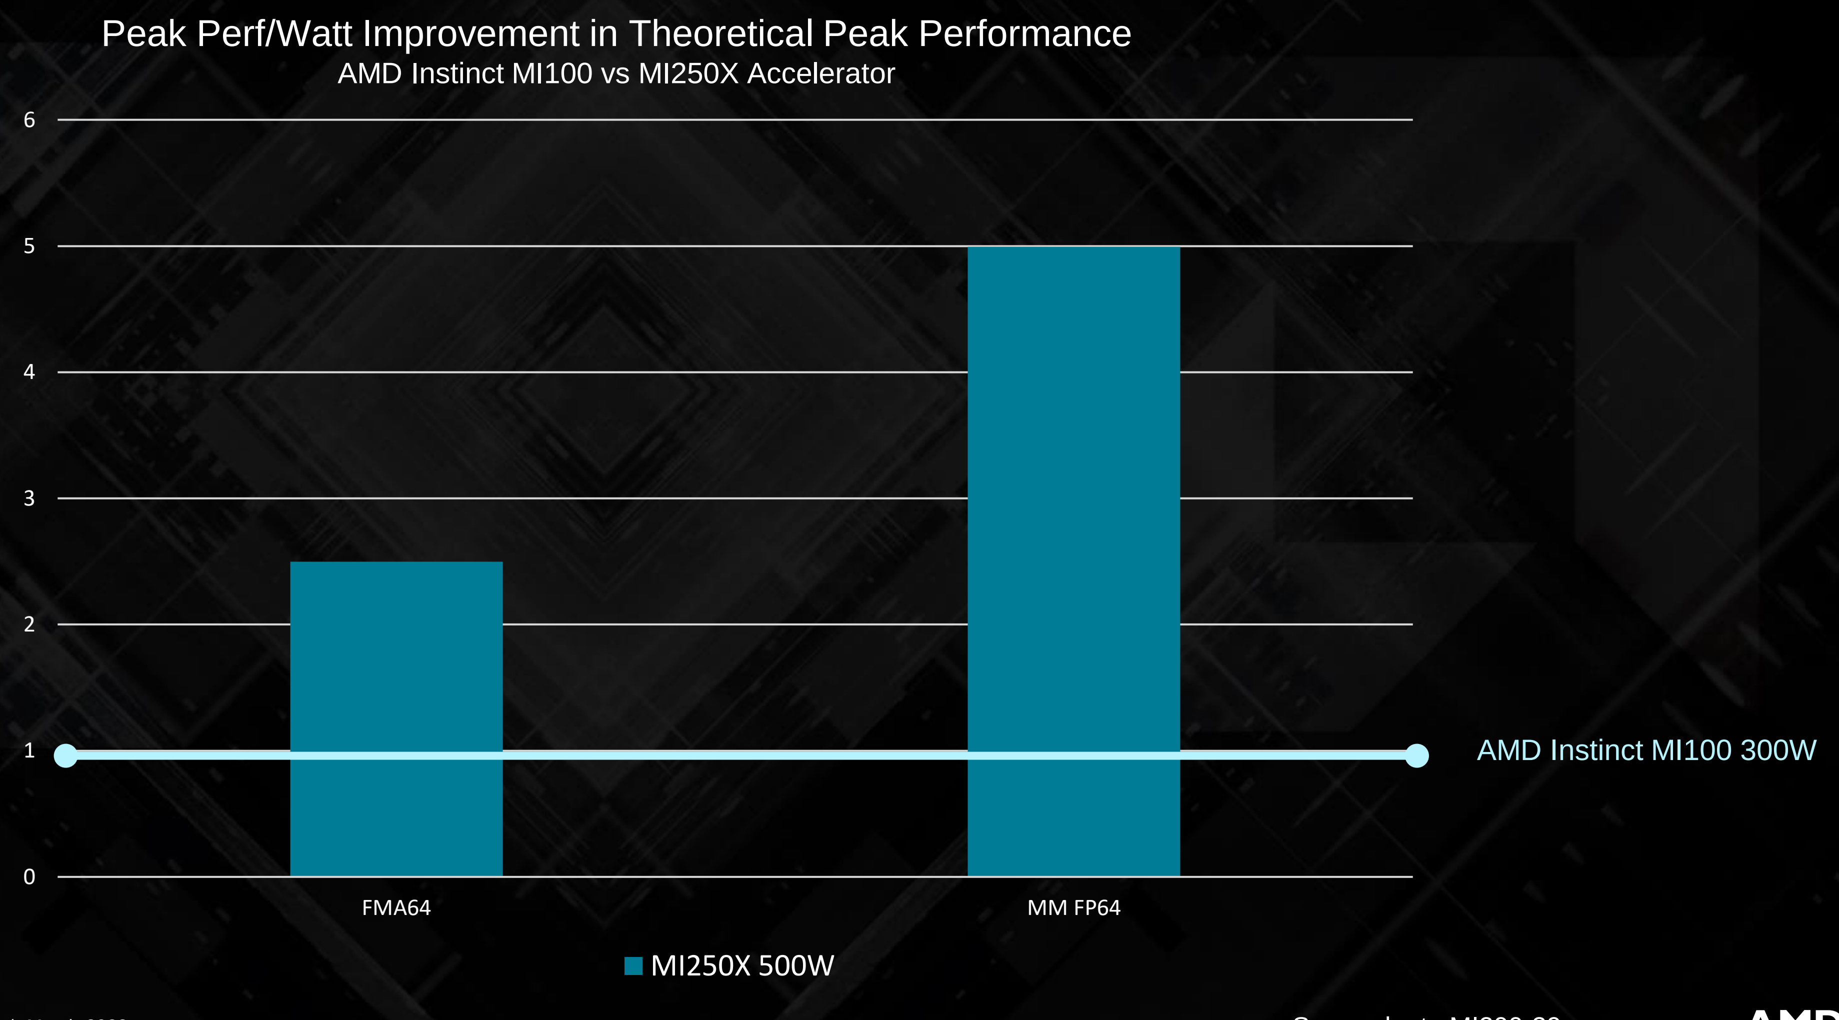

See endnote MI200-39

AMD Instinct MI200 Series
Up to 3.9X better Peak Perf/Watt Than Competition For HPC Workloads

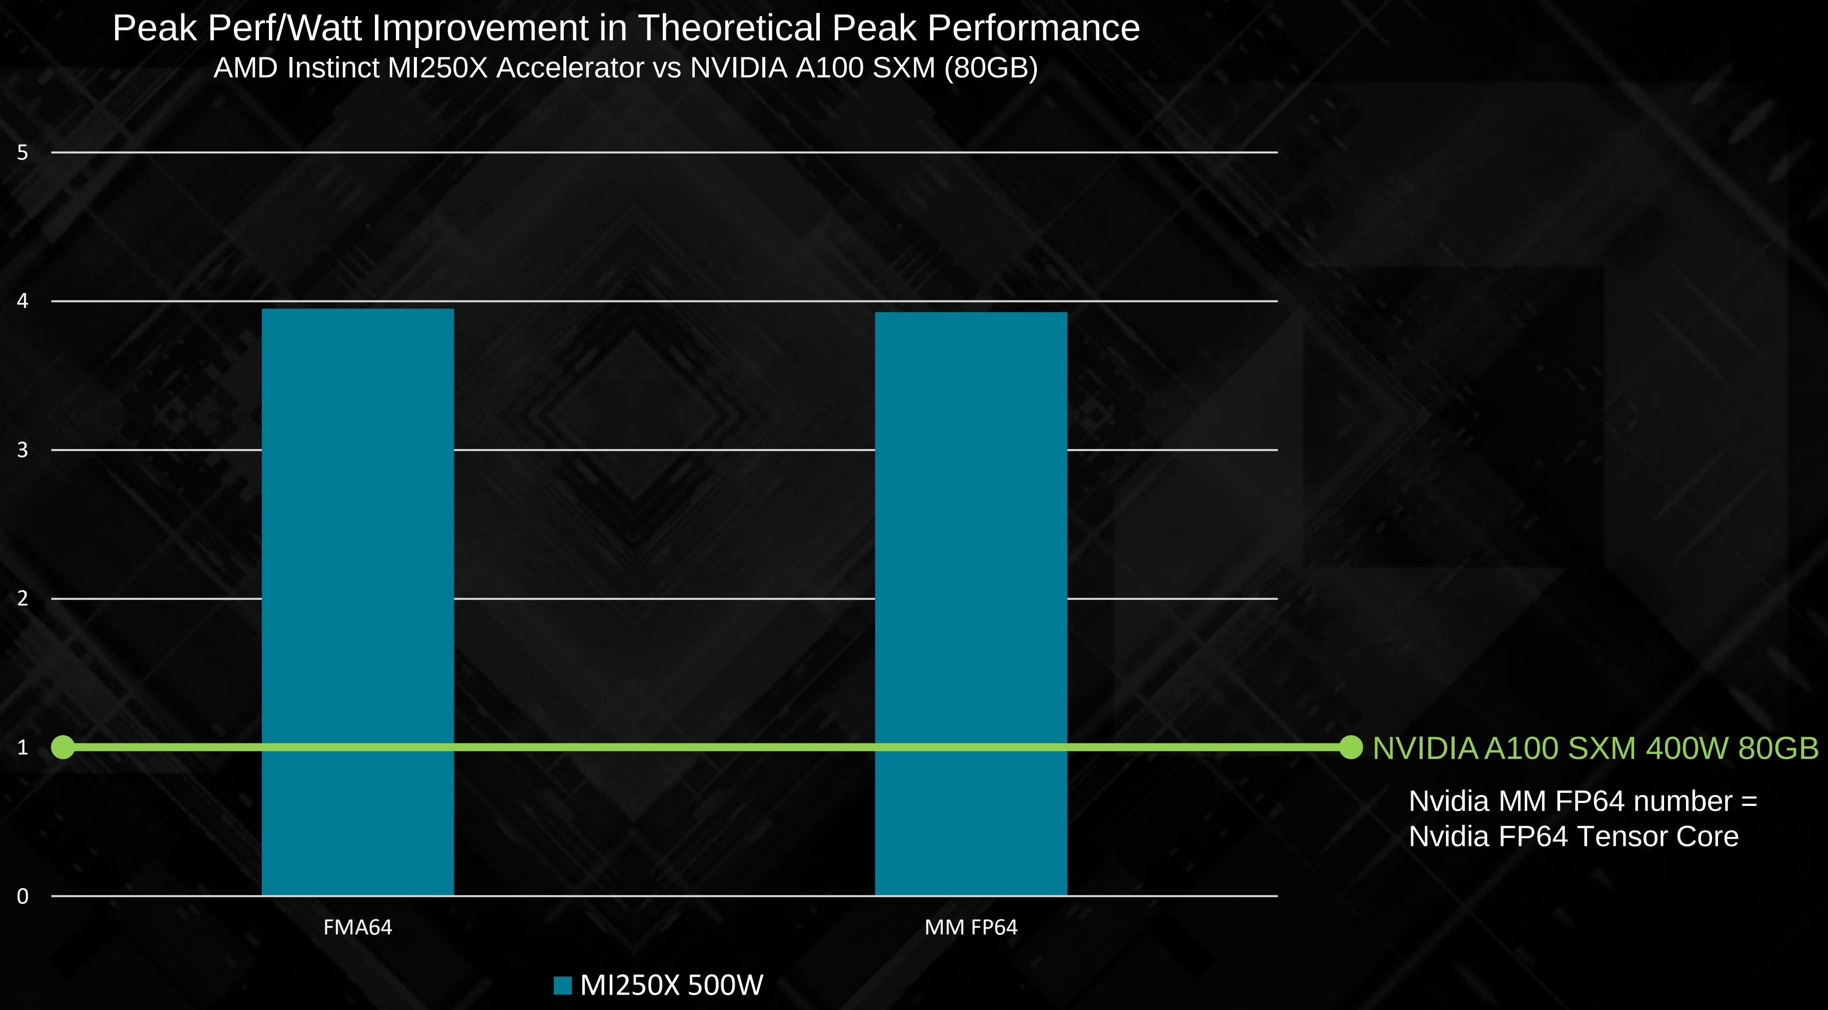

Footnote MI200-40

AMD

Ryzen 6000

Embargoed till 16 Apr

COMMERCIAL CLIENT INNOVATIONS

ADVANCING SUSTAINABILITY WITH PROCESSOR DESIGN

LEADERSHIP PERFORMANCE PER WATT

CUTTING-EDGE PROCESS TECHNOLOGY

ENHANCED POWER MANAGEMENT FEATURES

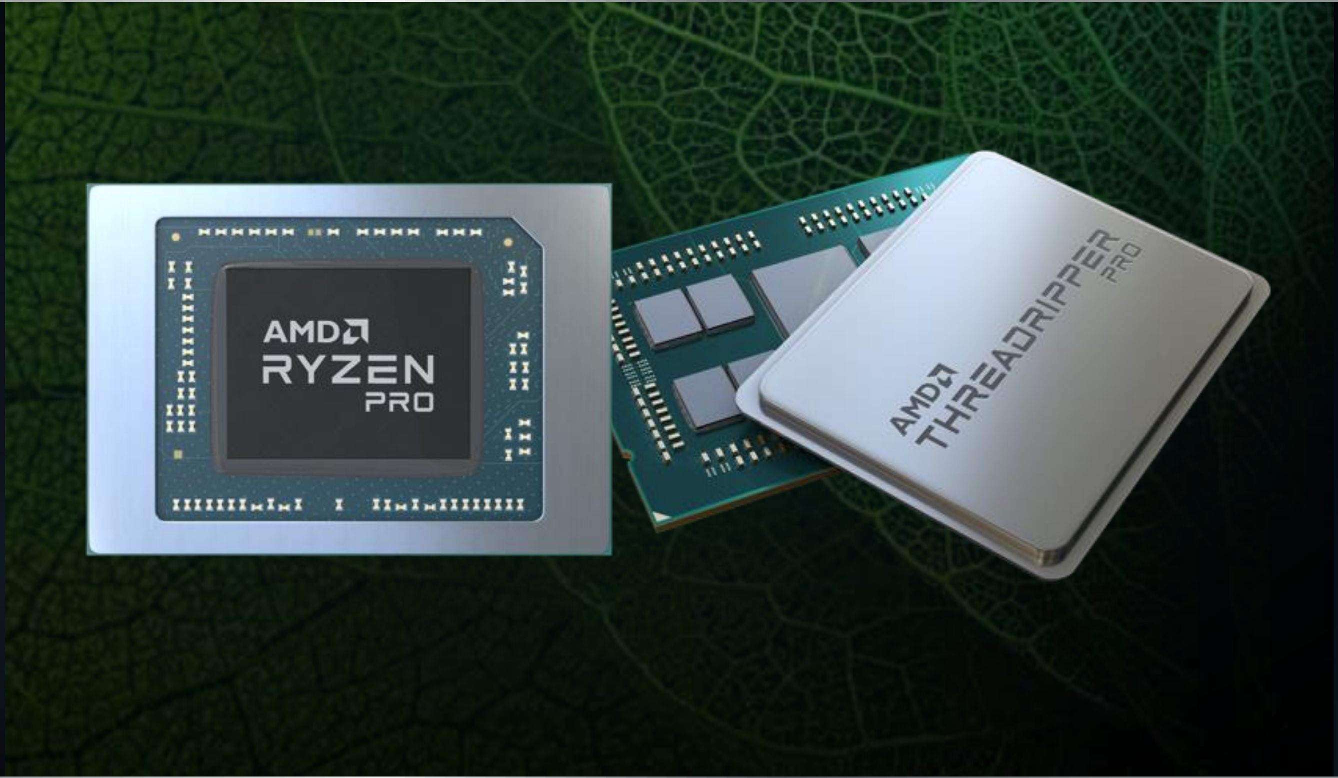

SUPPORT FOR LATEST LOWER POWER MEMORY

HIGH PERFORMANCE ULTRA-EFFICIENT

OPTIMIZED WORKLOADS

AMD RYZEN™ PROCESSORS ADVANCING SUSTAINABILITY

BUILDING A BETTER TOMORROW TODAY

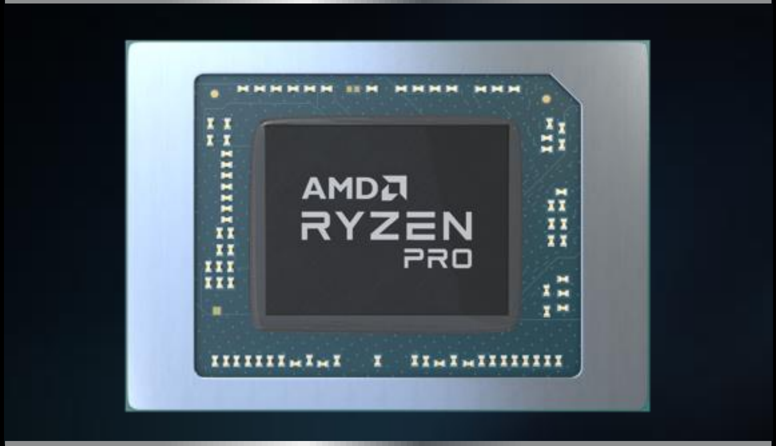

INDUSTRY LEADING POWER EFFICIENT CPUS

- 50% power improvement 2018 to 2022[1]
- 2.5X performance increase vs. 4-year-old laptop[2]
- AMD Exceeding Energy Star 8.0 requirements by 41% in 2022[3]

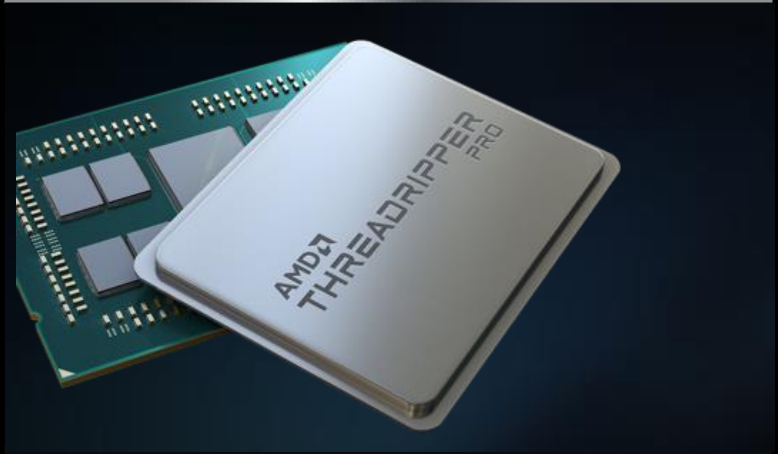

POWERFUL, EFFICIENT DESKTOP PROCESSING

- Yields up to 2x more performance per watt vs the competition
- Up to 67% less watts per core
- Up to 39% faster rendering

See endnote CGP-24

COLLABORATION WITH ENERGY REGULATIONS

- California Energy Commission (CEC)
- US EPA ENERGY STAR
- European Union Eco-design for computers (Lot 3) and servers (Lot 9)
- China National Institute of Standardization (CNIS

Not a comprehensive list

1: Based on APU power measurements by AMD labs using MobileMark 2014 (Ryzen 2500U vs Ryzen 5800U as of February 2022. 2: Based on measurements by AMD labs as of February 2022 of the Ryzen 2500U and Ryzen 5800U using a 50:50 weighted average to compare the increase in Cinebench 15nt and 3DMark performance. 3: Based on measurements by AMD labs as of February 2022 of the AMD Ryzen 7 Pro 5800U against Energy Star 8.0 requirements.

ESTIMATED SAVINGS OVER 4 YEARS ACROSS 10K LAPTOPS

UPDATING RYZEN 2000 SERIES TO RYZEN 5000 SERIES

272,000
Kilowatt Hours[1]

183 METRIC TONS
CO2 reduced[2]

3,187
Tree seedlings grown[2]

2.4X
Performance increase[3]

1Based on Energy Star measurements of Ryzen 2500 vs. Ryzen 5800 as measure in AMD lab. 2Estimate Based on Energy Saved by Energy Star calculations across 10K Units using EPA calculatorhttps://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator. 3Based on 50:50 weighted avg increase in Cinebench 15nt, 3DMark across Ryzen 2500 to Ryzen 5800.

AMD

Endnotes

1. Based on APU power measurements by AMD labs using MobileMark 2014 (Ryzen 2500U vs Ryzen 5800U) as of February 2022.
2. Based on measurements by AMD labs as of February 2022 of the Ryzen 2500U and Ryzen 5800U using a 50:50 weighted average to compare the increase in Cinebench 15nt and 3DMark performance.
3. Based on measurements by AMD labs as of February 2022 of the AMD Ryzen 7 Pro 5800U against Energy Star 8.0 requirements.
4. Includes AMD 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.

5 Based on 2015-2020 industry trends in energy efficiency gains and data center energy consumption in 2025.

6. Scenario based on all AI and HPC server nodes globally making similar gains to the AMD 30x goal, resulting in cumulative savings of up to 51.4 billion kilowatt-hours of electricity from 2021-2025 relative to baseline 2020 trends. Assumes $0.12 cents per kwh x 51.4 billion kwh = $6.2 million USD. Metric tonnes of CO2e emissions, and the equivalent estimate for wind turbines and tree plantings, is based on entering electricity savings into the U.S. EPA Greenhouse Gas Equivalency Calculator on 12/1/2021. https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator

CGP-24 Based on internal AMD analysis of benchmarks as of January 31st, 2022, evaluating the V-Ray rendering performance and TDP of an AMD Ryzen Threadripper Pro 5995WX reference system (280W) configured with 8x32GB DDR4, NVIDIA Quadro RTX A5000, 1TB SSD, Win 11 vs. a similarly configured BOXX APEXX4 workstation with TWO an Intel® Xeon® W-8280 server processors (410W). Results may vary. CGP-24

Endnotes

MLNTCO-020: 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 Bare Metal Server Greenhouse Gas Emissions TCO (total cost of ownership) Estimator Tool compares the selected AMD EPYC™ and Intel® Xeon® CPU based server solutions required to deliver a TOTAL_PERFORMANCE of 10000 units of integer performance based on the published scores for Intel Xeon and AMD EPYC CPU based servers. This estimation reflects a 3-year time frame.
This analysis compares a 2P AMD EPYC EPYC_7763 powered server with a SPECrate®2017_int_base score of 861, https://spec.org/cpu2017/results/res2021q4/cpu2017-20211121-30148.pdf; compared to a 2P Intel Xeon Platinum_8380 based server with a SPECrate®2017_int_base score of 602, https://spec.org/cpu2017/results/res2021q2/cpu2017-20210521-26364.pdf.

Both AMD EPYC and Intel based servers use the same cost for the following elements of the analysis: server chassis size of 2RU at a cost of $2500 per chassis; internal storage $380; physical servers managed per admin: 30; fully burdened cost per admin $110500; server rack size of 42; space allowance per rack of 27 sq feet; monthly cost of data center space $20 per sq foot; cost per kW for power $0.12; power drop per rack of 8kW; and a PUE (power usage effectiveness) of 1.7.

The EPYC powered solution is estimated to take: 12 total 2P EPYC_7763 powered servers at a hardware only acquisition cost of $23748 per server, which includes $7890 per CPU, total system memory of 1024GB, which is 8GB of memory / core and a total system memory cost of $5088; internal storage cost of $380. The total estimated AMD EPYC hardware acquisition cost for this solution is $284976. Each server draws ~755.1412kWhr per month. For the 3 years of this EPYC powered solution analysis the: total solution power cost is ~$66548.88 which includes the PUE factor; the total admin cost is ~$132600, and the total real estate cost is ~$38880, using 2 racks. The total 3 TCO estimate for the AMD solution is $523004.88.

The Intel based solution is estimated to take 17 total 2P Platinum_8380 powered servers at a hardware only acquisition cost of $24206 per server, which includes $8099 per CPU, total system memory of 1024GB, which is 12.8GB of memory / core and a total system memory cost of $5088; internal storage cost of $380. The total estimated Intel hardware acquisition cost for this solution is $411502. Each server draws ~751.4912kWhr per month. For the 3 years of this Intel based solution analysis the: total solution power cost is ~$93822.048 which includes the PUE factor; the total admin cost is ~$187851, and the total real estate cost is ~$58320 using 3 racks. The total 3 TCO estimate for the Intel solution is $751495.048.
AMD EPYC powered servers have a $228490 lower 3-year TCO.

Delivering 10000 estimated score of SPECrate®2017_int_base performance produces the following estimated results: the AMD EPYC solution requires 29% fewer servers [1-(AMD server count / Intel server count)]; 33% less space [1-(AMD rack count / Intel rack count)]; 29% less power [1-(AMD power cost / Intel power cost)]; providing a 30% lower 3-year TCO [1-(AMD TCO / Intel TCO)]. delivering ~98 or ~1% Better w/ AMD SPECrate®2017_int_base solution score
AMD EPYC_7763 powered servers save ~227276.4kWh of electricity for the 3 years of this analysis. 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', the AMD EPYC powered server saves ~103.01 Metric Tons of CO2 equivalents. This results in the following estimated savings based on United States data,
Emissions Avoided equivalent to one of the following:
- 22 USA Passenger Cars Not Driven for 1 year; or; or
- 7.45 USA Passenger Cars Not Driven Annually; or; or
- 11640 Gallons of Gasoline Not Used; or; or

or Carbon Sequestered equivalent to:
- 1700 Tree Seedlings Grown for 10 years in USA; or; or
- 41.2 Acres of USA Forests Annually.

The 2020 Grid Electricity Emissions Factors v1.4 – September 2020 data used in this analysis can be found at https://www.carbonfootprint.com/docs/2020_09_emissions_factors_sources_for_2020_electricity_v14.pdf and the US EPA Greenhouse Gas Equivalencies Calculator used in this analysis can be found athttps://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator

AMD processor pricing based on 1KU price as of Sept 2021. Intel® Xeon® Scalable Gen 1 and Gen 2 CPU data and pricing from https://ark.intel.com as of September 2021. Intel Xeon Gen3 Scalable Ice Lake pricing and data from https://newsroom.intel.com/wp-content/uploads/sites/11/2021/05/3rd-Gen-Intel-Xeon-Scalable-Processor-SKU-Stack-with-RCP.pdf on 09/01/2021. All pricing is in USD.

SPECrate® scores as of Jan 14, 2022. SPEC®, SPECrate® and SPEC CPU® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information.
AMD EPYC performance numbers based on the identified benchmark reported scores or the user provided score where indicated. Product and company names are for informational purposes only and may be trademarks of their respective owners.
Results generated by: AMD EPYC™ BARE METAL SERVER and GREENHOUSE GAS EMISSIONS TCO ESTIMATION TOOL; VERSION 4.2

Endnotes

MLNTCO-008a: 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 Virtualization TCO (total cost of ownership) Estimator tool compares the 2P AMD EPYC™ and the 2P Intel® Xeon® server solutions required to deliver 320 total virtual machines (VM), requiring 1 core and 8GB of memory per VM. The analysis includes both hardware and virtualization software components. Hardware costs (CPU + memory + storage + chassis): The 2P AMD 64 core EPYC_7713 processor used in this solution analysis provides 128 total cores per server, each processor cost $7060 and the server uses 32 x 32GB DIMMs to achieve the minimum required memory footprint, in a 1RU server chassis that cost $2200, and requires 1 server racks. The AMD solution has a total estimated hardware acquisition cost of $66612. The 40 core Intel Xeon Platinum_8380 processor used in this solution analysis provides 80 total cores per server. Each processor cost $8099 and the server uses 24 x 32GB DIMMs to achieve the minimum required memory footprint, in a 2RU server chassis that cost $2500 and requires 1 server racks. The Intel solution has a total estimated hardware acquisition cost of $92824.
PROCESSOR COSTS: AMD processor pricing based on 1KU price as of March 2021. Intel® Xeon® Scalable processor data and pricing from https://ark.intel.com as of September 2020.
OPERATING COSTS: AMD has estimated OpEx costs as follows: a hardware admin cost of $33150, a real estate cost of $19440, and a power cost of $17905.453056, for a total estimated 3 year TCO cost (hardware cost and operating expense) of $137107 with AMD. Estimated OpEx costs for Intel are: hardware admin cost of $44200, real estate cost of $19440, and power cost of $23243, for a total estimated 3 year TCO cost (hardware cost and operating expense) of $179707, with Intel. AMD has an estimated 24% lower hardware TCO for this virtualization solution, 1 - ($137107 ÷ $179707) = 24%. The core assumptions for this analysis are as follows: Cost of power @ $0.12 with kwatts (kW) of power to each rack and a PUE (power usage effectiveness) of 2 and a server rack size of 42RU. Each server has 1 hard drives drawing 3 watts each. Server Admin annual salary is $85000 managing 30 physical servers with a salary burden rate of 30%. The VM Admin salary is $85000, with a burden rate of 30% and managing 400 VMs. Each server uses the following power: AMD EPYC_7713 uses 959.44 W/hr; the Intel Platinum_8380 uses 934.08 watts per hour.
VIRTUALIZATION TCO: Analysis is based on the following estimiates: 3 year Virtualization (hardware, operating, and software cost) for the Intel solution is $540401 and $473927 for the AMD solution. This means that the AMD solution is ~12% less expensive over three years. 1 - ($473927 ÷ $540401) = 12%. The AMD solution 1st year TCO per VM of $781.66 where the Intel 1st yr. solution is $955.25. The AMD 1st year TCO per VM is $173.59, or ~18% lower than Intel. The virtualization software used in this analysis is VMware with a VMware® vSphere Enterprise Plus w/ Production support license. This analysis uses license pricing of $5968.36 per Socket + Core with 3 year support. More information on VMware software can be found @ https://store-us.vmware.com/vmware-vsphere-enterprise-plus-284281000.html.
For 320 VMs with 1 core(s) per VM, and 8 GB of memory per VM, the Intel Platinum_8380 processor requires 4 servers, and 16 licenses. The AMD EPYC_7713 solution requires 3 servers and 12 licenses. The AMD solution requires 25% fewer servers than the Intel solution.
The AMD server and virtualization license cost are $138232.32, and the Intel cost are $188317.76. Hardware and virtualization cost are ~$50085 or ~27% Lower w/ AMD. Virtualization software pricing as of 09/14/2021. Third party names are for informational purposes only and may be trademarks of their respective owners. All pricing is in USD.
Results generated by: AMD EPYC™ SERVER VIRTUALIZATION TCO ESTIMATION TOOL VERSION: v9.43

Endnotes

EPYC-028: As of 2/2/22, of SPECpower_ssj® 2008 results published on SPEC's website, the 55 publications with the highest overall efficiency results were all powered by AMD EPYC processors. More information about SPEC® is available at http://www.spec.org. SPEC and SPECpower are registered trademarks of the Standard Performance Evaluation Corporation.
Links to these 55 results are:

1 http://www.spec.org/power_ssj2008/results/res2020q4/power_ssj2008-20200918-01047.html
2 http://www.spec.org/power_ssj2008/results/res2020q4/power_ssj2008-20200918-01046.html
3 http://www.spec.org/power_ssj2008/results/res2021q2/power_ssj2008-20210324-01091.html
4 http://www.spec.org/power_ssj2008/results/res2020q2/power_ssj2008-20200519-01031.html
5 http://www.spec.org/power_ssj2008/results/res2021q1/power_ssj2008-20210309-01077.html
6 http://www.spec.org/power_ssj2008/results/res2020q2/power_ssj2008-20200407-01022.html
7 http://www.spec.org/power_ssj2008/results/res2021q2/power_ssj2008-20210408-01094.html
8 http://www.spec.org/power_ssj2008/results/res2020q2/power_ssj2008-20200519-01034.html
9 http://www.spec.org/power_ssj2008/results/res2021q2/power_ssj2008-20210413-01095.html
10 http://www.spec.org/power_ssj2008/results/res2021q1/power_ssj2008-20210309-01078.html
11 http://www.spec.org/power_ssj2008/results/res2020q2/power_ssj2008-20200519-01032.html
12 http://www.spec.org/power_ssj2008/results/res2020q2/power_ssj2008-20200407-01023.html
13 http://www.spec.org/power_ssj2008/results/res2020q2/power_ssj2008-20200407-01025.html
14 http://www.spec.org/power_ssj2008/results/res2020q2/power_ssj2008-20200519-01033.html
15 http://www.spec.org/power_ssj2008/results/res2020q2/power_ssj2008-20200407-01024.html
16 http://www.spec.org/power_ssj2008/results/res2021q4/power_ssj2008-20211001-01130.html
17 http://www.spec.org/power_ssj2008/results/res2021q2/power_ssj2008-20210602-01106.html
18 http://www.spec.org/power_ssj2008/results/res2021q2/power_ssj2008-20210602-01105.html
19 http://www.spec.org/power_ssj2008/results/res2020q3/power_ssj2008-20200714-01039.html
20 http://www.spec.org/power_ssj2008/results/res2020q1/power_ssj2008-20191125-01012.html
21 http://www.spec.org/power_ssj2008/results/res2021q2/power_ssj2008-20210615-01111.html
22 http://www.spec.org/power_ssj2008/results/res2020q3/power_ssj2008-20200714-01040.html
23 http://www.spec.org/power_ssj2008/results/res2020q2/power_ssj2008-20200324-01021.html
24 http://www.spec.org/power_ssj2008/results/res2020q1/power_ssj2008-20191125-01011.html
25 http://www.spec.org/power_ssj2008/results/res2020q2/power_ssj2008-20200313-01020.html
26 http://www.spec.org/power_ssj2008/results/res2020q2/power_ssj2008-20200313-01019.html
27 http://www.spec.org/power_ssj2008/results/res2020q1/power_ssj2008-20200310-01018.html
28 http://www.spec.org/power_ssj2008/results/res2019q3/power_ssj2008-20190717-00987.html
29 http://www.spec.org/power_ssj2008/results/res2019q3/power_ssj2008-20190717-00988.html
30 http://www.spec.org/power_ssj2008/results/res2019q3/power_ssj2008-20190909-01004.html
31 http://www.spec.org/power_ssj2008/results/res2019q3/power_ssj2008-20190717-00986.html
32 http://www.spec.org/power_ssj2008/results/res2021q1/power_ssj2008-20210221-01066.html
33 http://www.spec.org/power_ssj2008/results/res2019q3/power_ssj2008-20190717-00990.html
34 http://www.spec.org/power_ssj2008/results/res2019q3/power_ssj2008-20190717-00985.html
35 http://www.spec.org/power_ssj2008/results/res2020q3/power_ssj2008-20200728-01041.html
36 http://www.spec.org/power_ssj2008/results/res2021q1/power_ssj2008-20210221-01063.html
37 http://www.spec.org/power_ssj2008/results/res2019q3/power_ssj2008-20190716-00980.html
38 http://www.spec.org/power_ssj2008/results/res2021q1/power_ssj2008-20210221-01064.html
39 http://www.spec.org/power_ssj2008/results/res2021q1/power_ssj2008-20210221-01065.html
40 http://www.spec.org/power_ssj2008/results/res2019q3/power_ssj2008-20190716-00982.html
41 http://www.spec.org/power_ssj2008/results/res2021q1/power_ssj2008-20210223-01073.html
42 http://www.spec.org/power_ssj2008/results/res2020q2/power_ssj2008-20200407-01029.html
43 http://www.spec.org/power_ssj2008/results/res2020q2/power_ssj2008-20200407-01028.html
44 http://www.spec.org/power_ssj2008/results/res2019q3/power_ssj2008-20190716-00981.html
45 http://www.spec.org/power_ssj2008/results/res2019q4/power_ssj2008-20191203-01015.html
46 http://www.spec.org/power_ssj2008/results/res2021q1/power_ssj2008-20210222-01068.html
47 http://www.spec.org/power_ssj2008/results/res2020q2/power_ssj2008-20200407-01026.html
48 http://www.spec.org/power_ssj2008/results/res2021q1/power_ssj2008-20210223-01074.html
49 http://www.spec.org/power_ssj2008/results/res2019q3/power_ssj2008-20190911-01005.html
50 http://www.spec.org/power_ssj2008/results/res2021q1/power_ssj2008-20210222-01069.html
51 http://www.spec.org/power_ssj2008/results/res2019q3/power_ssj2008-20190730-00994.html
52 http://www.spec.org/power_ssj2008/results/res2021q1/power_ssj2008-20210222-01071.html
53 http://www.spec.org/power_ssj2008/results/res2020q2/power_ssj2008-20200407-01027.html
54 http://www.spec.org/power_ssj2008/results/res2019q3/power_ssj2008-20190717-00984.html
55 http://www.spec.org/power_ssj2008/results/res2021q1/power_ssj2008-20210222-01072.html

MI200-39
Calculations conducted by AMD Performance Labs as of Sep 15, 2021, for the AMD Instinct™ MI250X (128GB HBM2e OAM module) 500 Watt accelerator at 1,700 MHz peak boost engine clock resulted in 95.7 TFLOPS peak theoretical double precision (FP64 Matrix), 47.9 TFLOPS peak theoretical double precision (FP64 vector) floating-point performance. The AMD Instinct MI100 (32GB HBM2 PCIe® card) accelerator at 1,502 MHz peak boost engine clock resulted in 11.5 TFLOPS peak theoretical double precision (FP64 Matrix), 11.5 TFLOPS peak theoretical double precision (FP64 vector) floating point performance.

MI200-40
Calculations conducted by AMD Performance Labs as of Sep 15, 2021, for the AMD Instinct™ MI250X (128GB HBM2e OAM module) 500 Watt accelerator at 1,700 MHz peak boost engine clock resulted in 95.7 TFLOPS peak theoretical double precision (FP64 Matrix), 47.9 TFLOPS peak theoretical double precision (FP64 vector) floating-point performance. The Nvidia A100 SXM (80 GB) accelerator (400W) with boost engine clock of 1410 MHz results in 19.5 TFLOPS peak theoretical double precision (FP64 Tensor Core), 9.7TF TFLOPS peak theoretical double precision (FP64 vector) floating-point performance.