BUILDING AN AI READY ENTERPRISE

AMD
together we advance_

AMD

TABLE OF CONTENTS

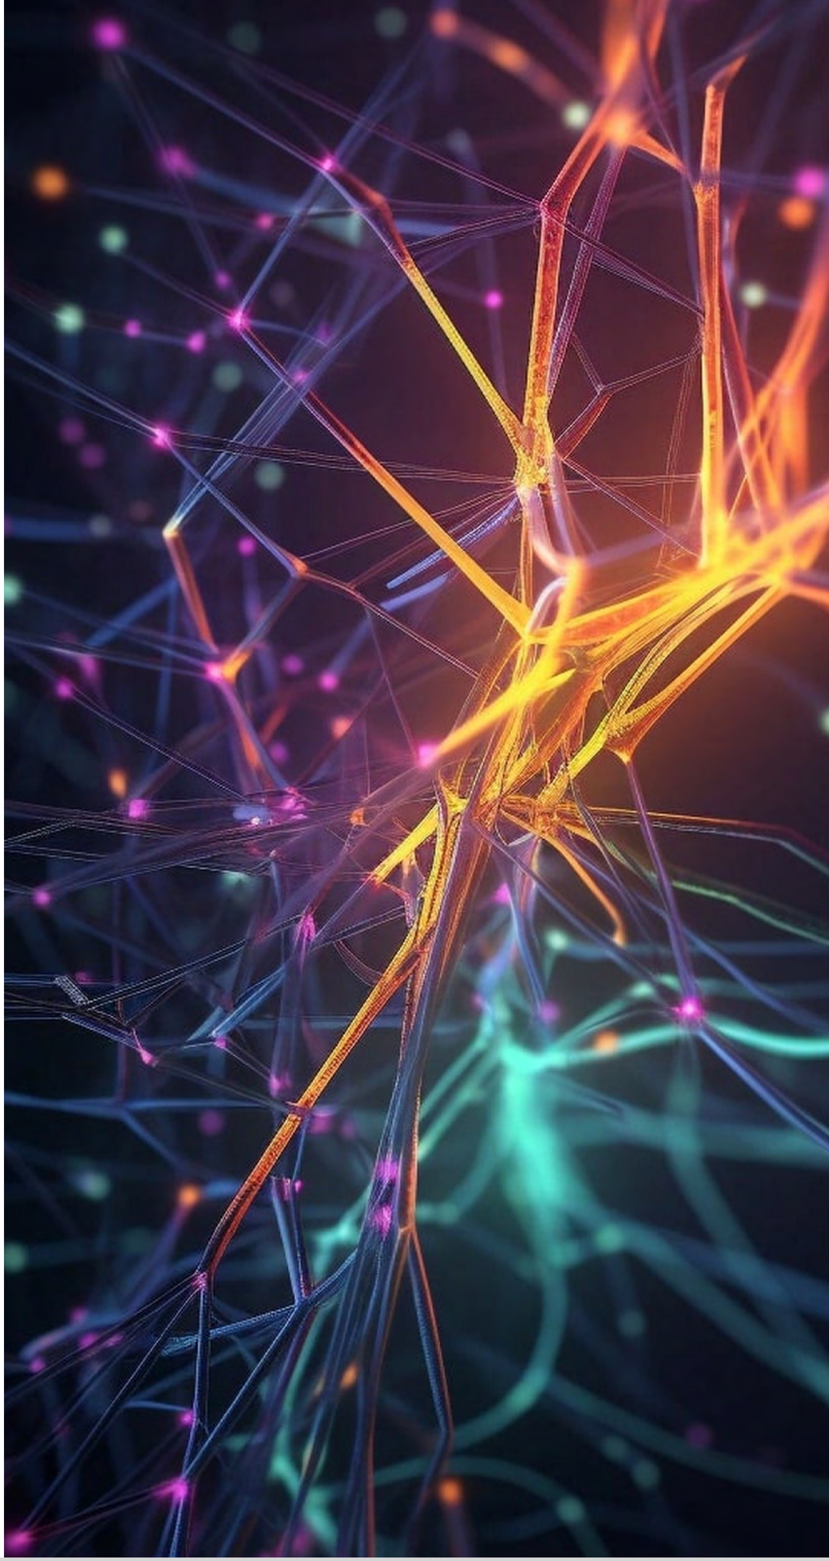

INTRODUCTION

THE ENTERPRISE AI IMPERATIVE

There's little historical precedent for the rapid growth of AI, which has swept rapidly through the enterprise technology landscape. AI is now a business imperative, with virtually every core function standing to benefit from faster decisions, increased productivity, and decreased operational costs.

But the rapid rise of AI is not a simple or seamless process for any organization. The challenges range from meeting new data center compute and energy demands, to aligning investments in resources with objectives around productivity gains or accelerated innovations.

This eBook is about navigating those challenges. Its insights will guide organizations on the necessary infrastructure modernizations that minimize the risk of costly miscues, missed innovation opportunities, and rising operational and maintenance costs. And it will chart a path to achieving accelerated business outcomes and sustained success.

3 KEY CHALLENGES IN ENTERPRISE AI ADOPTION

#1 LEGACY INFRASTRUCTURE ISN'T AI READY

The limitations of legacy infrastructure are the most pressing obstacle to enterprise adoption and scaling of AI.

Enterprises urgently need to accommodate AI workloads, and to do so rapidly and cost-effectively. Yet these workloads demand significant compute resources. They can't simply be layered onto already crowded data center infrastructures, where energy demands are a constant concern and CPU resources, storage and network bandwidth are operating at, or near, full capacity.

Adding additional racks of AI-capable hardware typically demands more density and floor space than legacy data centers were designed to accommodate. And making the attendant upgrades to power or cooling capacity is also likely to be prohibitively expensive. Modernization paths are further complicated by the risks involved with re-platforming business-critical legacy applications in parallel with building a new AI-ready environment.

Legacy infrastructure also carries high opportunity costs – resources spent on upkeep can't be redirected to AI innovation.

AMD

#2 AI IS INCREASING DEMAND FOR CONFIDENTIAL COMPUTING

As AI becomes a 'killer application' that touches every facet of business and personal data, trust in data confidentiality and integrity becomes essential.

Achieving high-performance AI often relies on heterogeneous hardware (CPUs, GPUs, and specialized accelerators), spread across multiple nodes and even multiple sites. Ensuring a secure, trusted boundary across all devices and network links is nontrivial. Even if data is encrypted at rest, vulnerabilities in virtualized or containerized environments may allow malicious hypervisors to access sensitive information.

In response, enterprises must integrate confidential computing features, such as Secure Encrypted Virtualization (SEV), into their AI workflows. And those protections must scale to tens or hundreds of GPU/accelerator nodes, combining with other data security techniques to create a trusted AI execution model.

#3 IT'S DIFFICULT TO INVEST DECISIVELY WHILE STAYING FLEXIBLE

The objectives of AI initiatives can broadly be grouped into two categories: productivity enhancements or innovation. Each one places different demands on data center infrastructure.

AI initiatives focused on increasing productivity through automation or improving workflow efficiency might primarily be characterized as inference workloads with lighter training requirements. They can often be comfortably handled by the latest generation of CPUs. On the other hand, projects targeting entirely new lines of revenue or breakthrough insights will have heavier training requirements that might necessitate investment in new GPU resources.

Balancing the portfolio of AI investments between these two objectives is a core strategic challenge. Advances in AI frameworks and hardware are arriving quickly, with each new generation of hardware or machine learning framework forcing an update to cost-performance calculations.

KEY INFRASTRUCTURE AREAS FOR AI MODERNIZATION

Accelerating business outcomes with AI means addressing each of these challenges through a series of targeted data center modernizations.

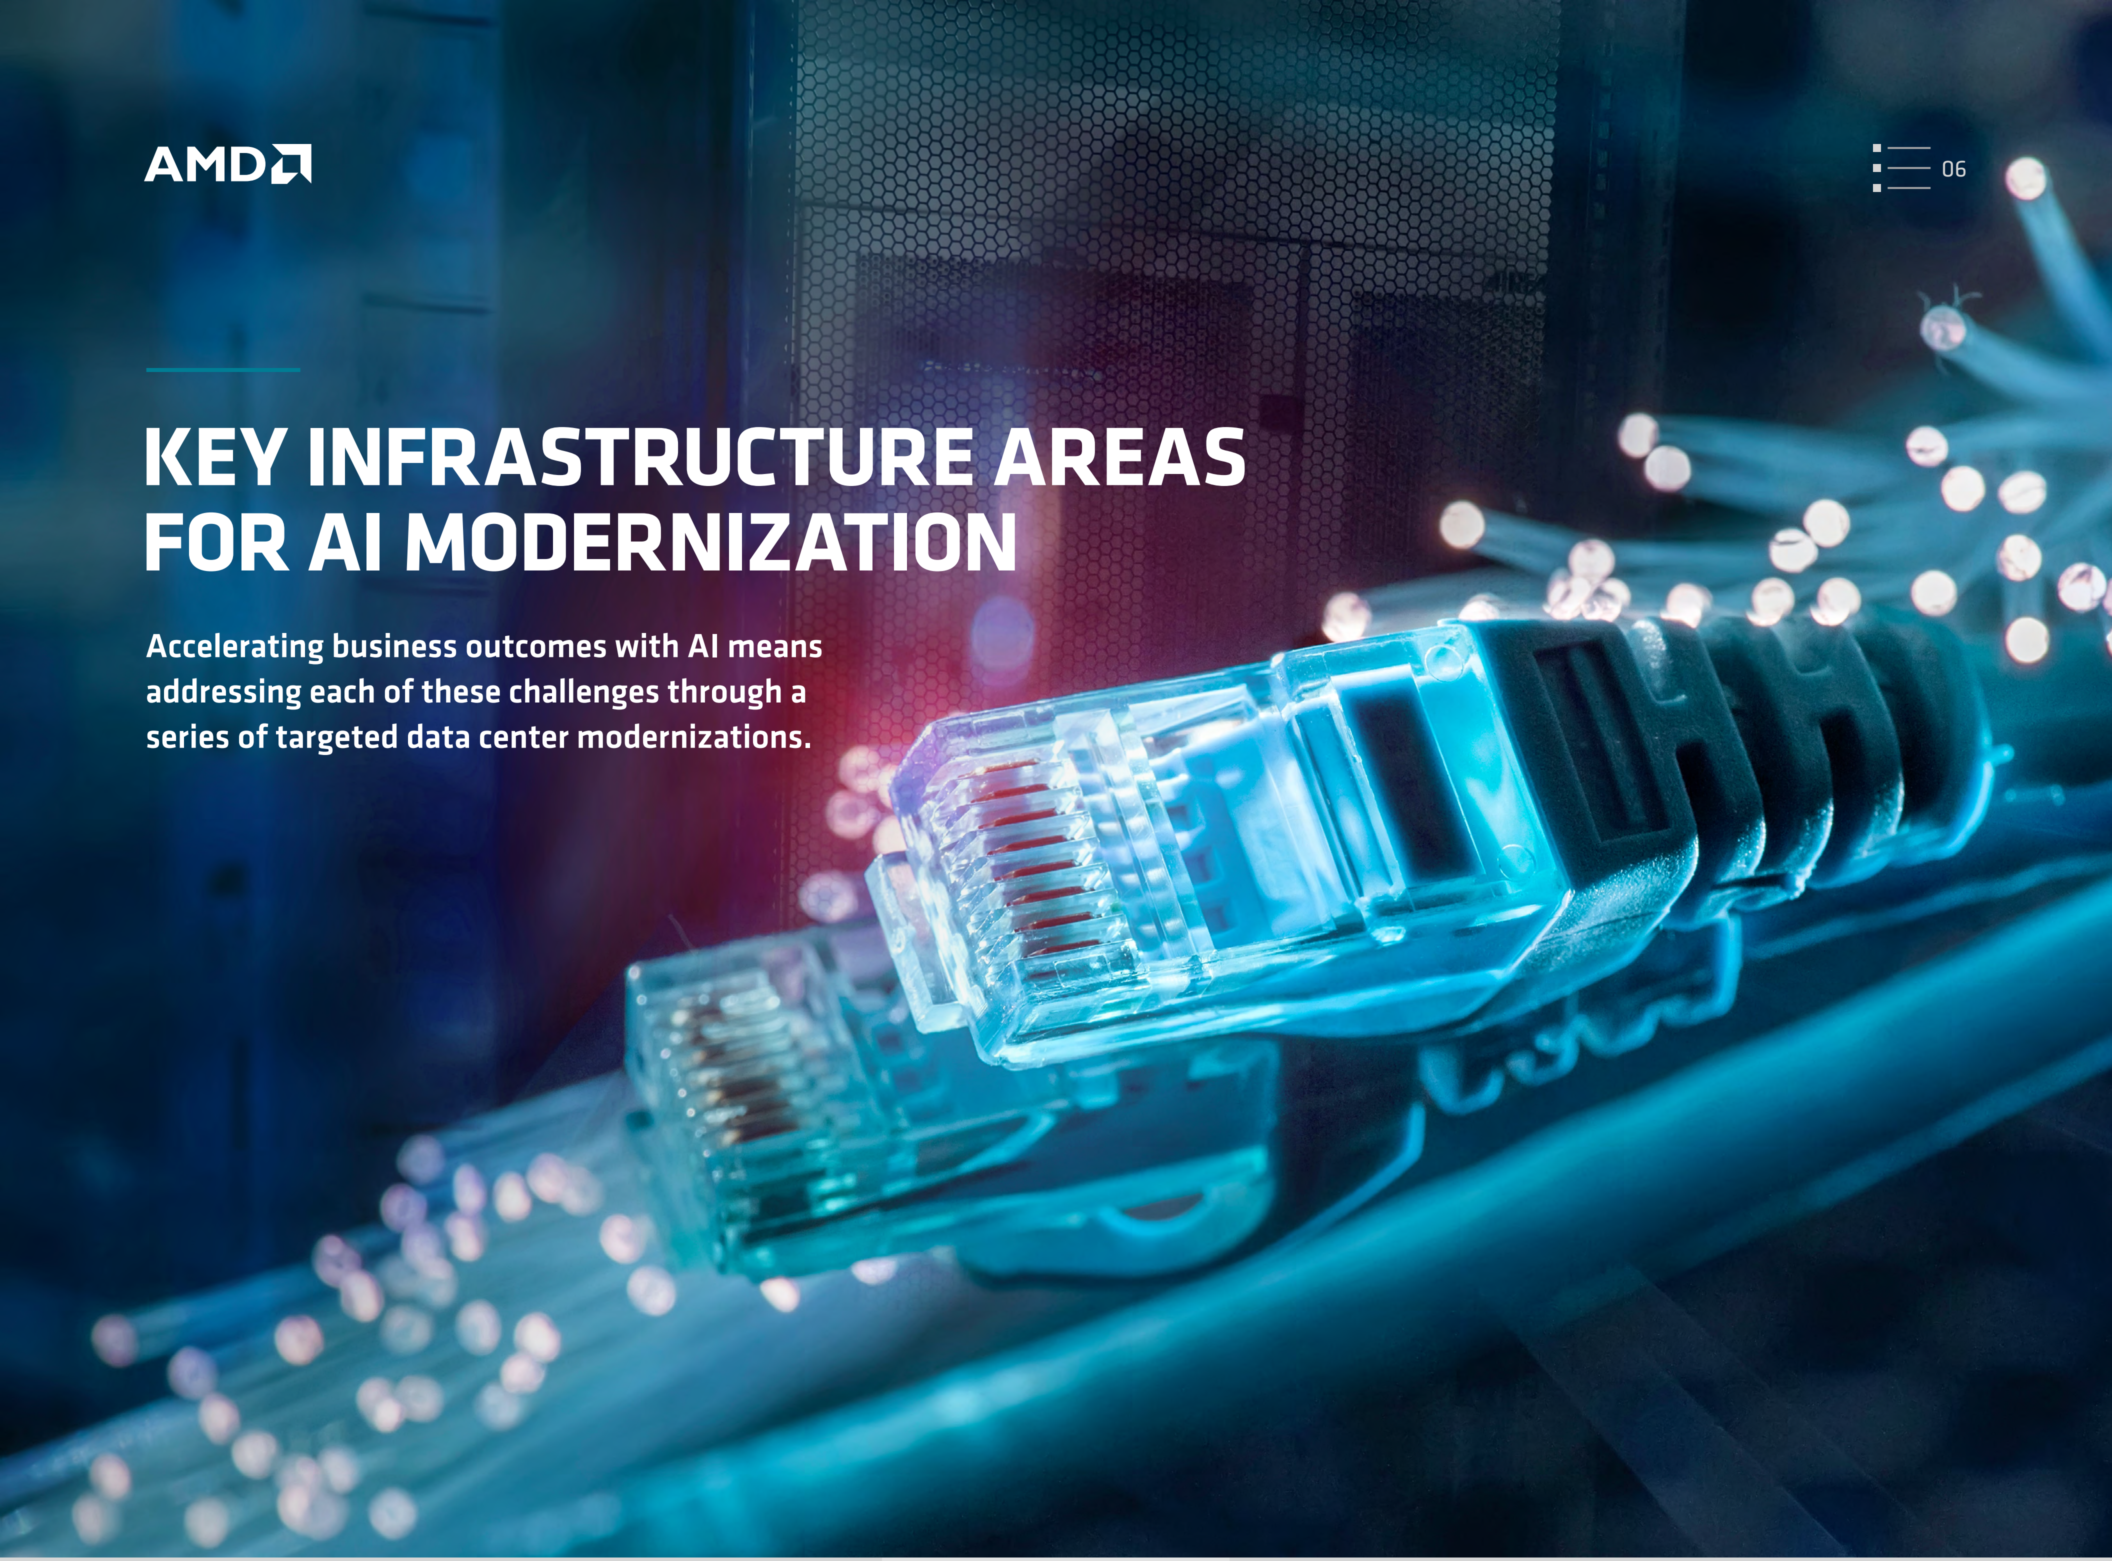

AMD

KEY INFRASTRUCTURE AREAS FOR AI MODERNIZATION

OPEN THE DOOR TO AI EXPANSION WITH SERVER CONSOLIDATION

Server consolidation drives space and energy savings in the data center, and throws open the doors to AI expansion.

Modern CPUs now offer dozens – sometimes hundreds – of cores per socket. This enables massive parallelism for AI data pre-processing and small- to medium-scale inference tasks, and means one server can replace up to seven legacy machines.[1] That opens up floor space for additional AI-centric racks, if required, and their attendant advanced cooling infrastructure.

The next generation of CPU architectures also offer greater performance per-watt compared to previous generations, which drives lower power utilization and reduces the ongoing operational costs of running AI workloads.

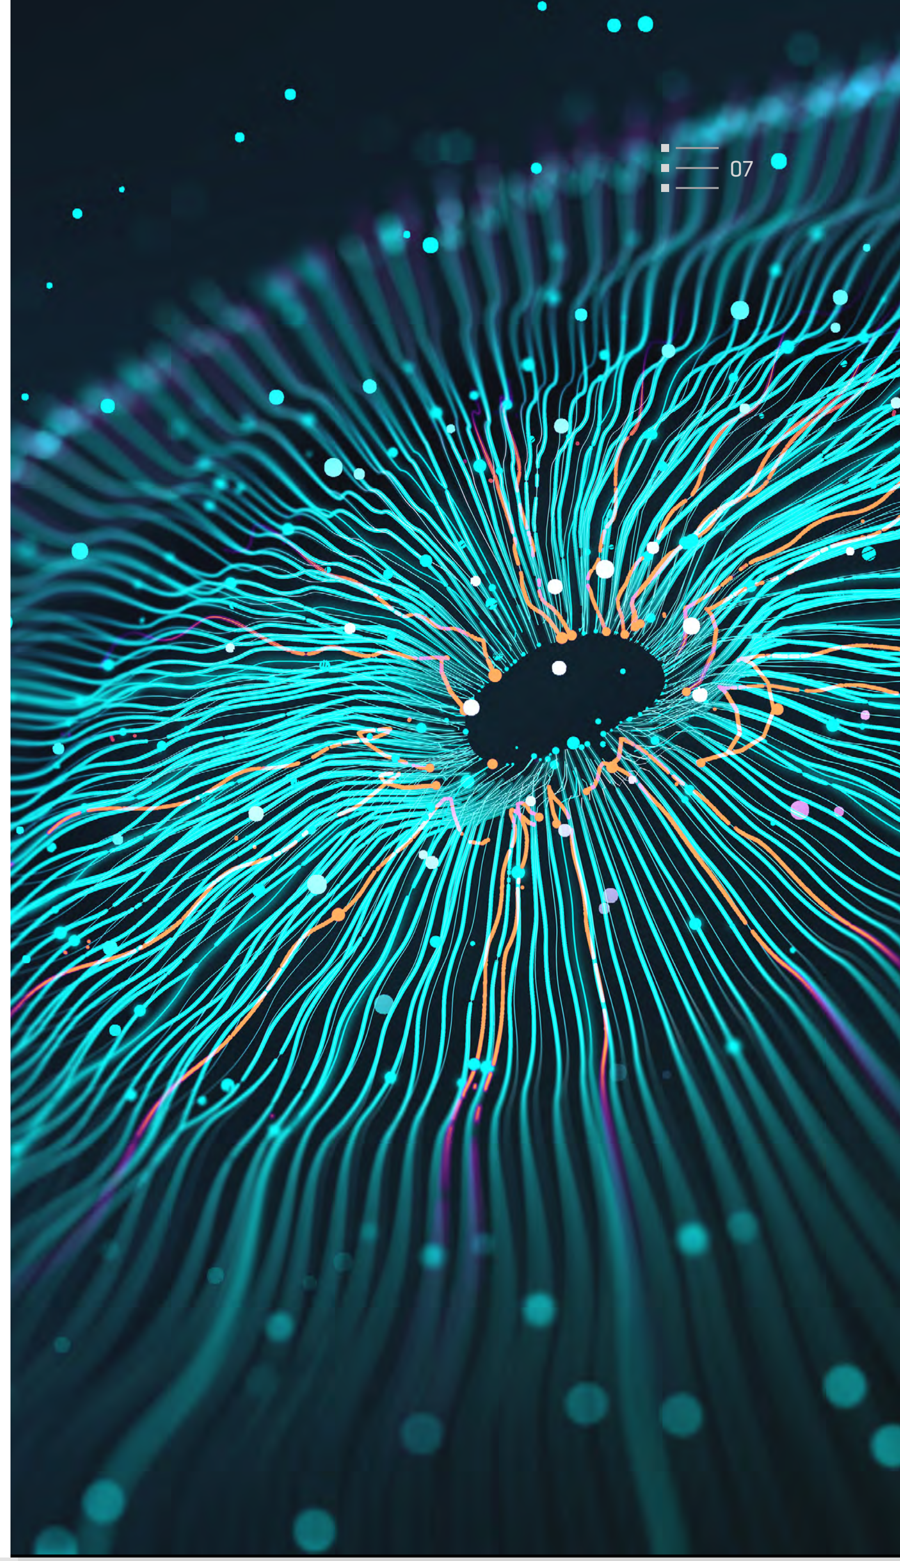

AMD EPYC™ PROCESSORS: THE LEADING CPU FOR AI[2]

AMD EPYC™ 9005 Series processors can match integer performance of legacy hardware with up to 86% fewer racks,[3] delivering up to 3x the Machine Learning throughput of 64-core Intel Xeon 8592+ processors.[4]

Hardware-level encryptions, such as Secure Encrypted Virtualization (SEV), also mean models and data can remain encrypted during training or inference, so data remains confidential even in multi-tenant environments.

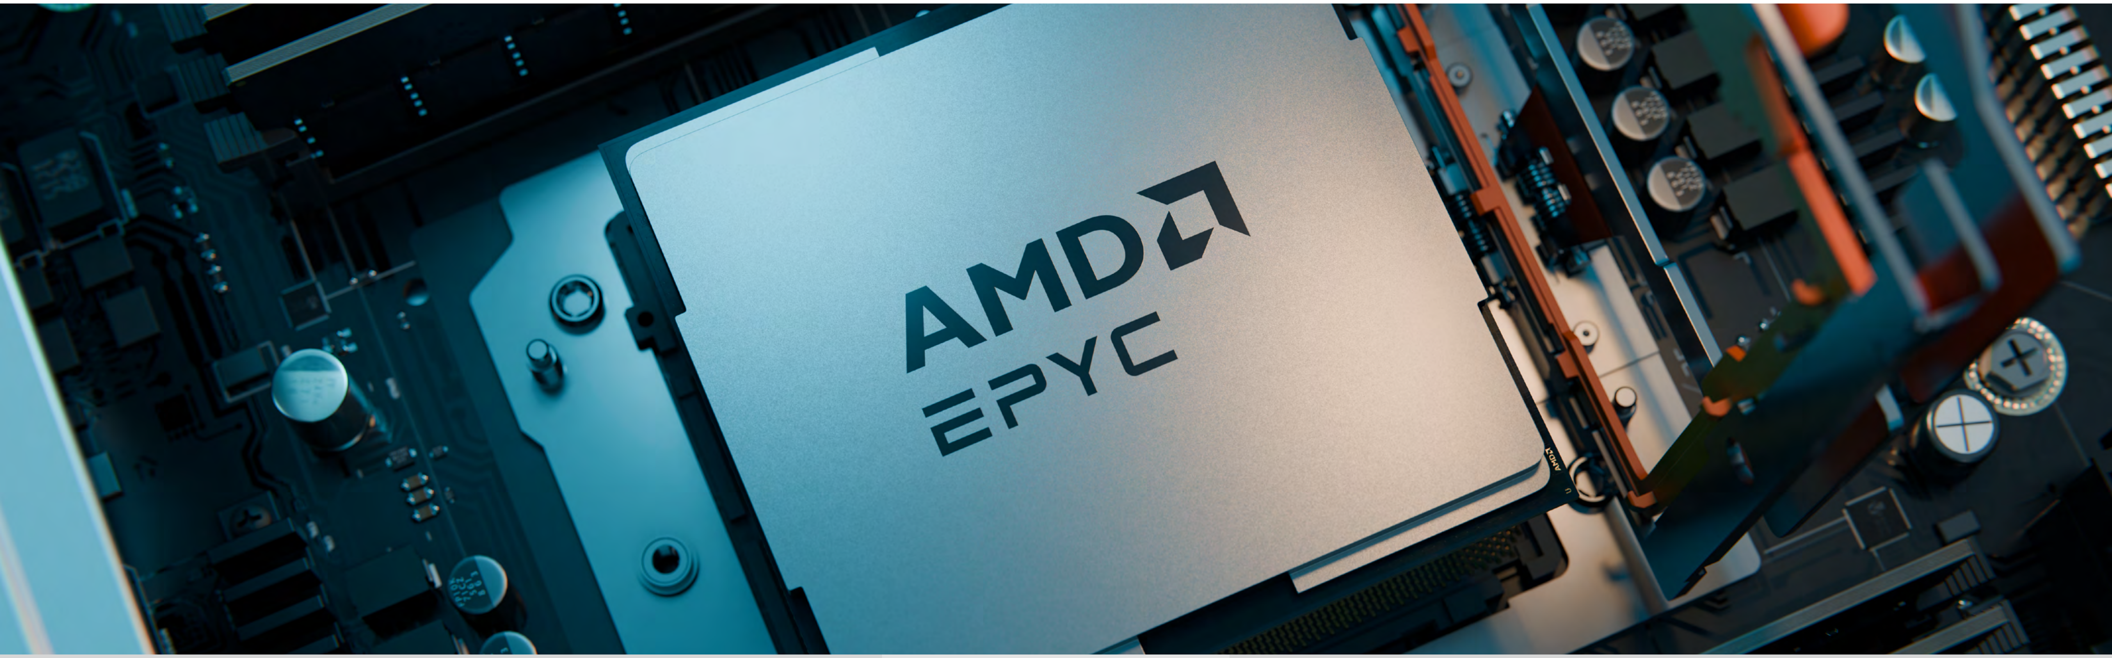

KEY INFRASTRUCTURE AREAS FOR AI MODERNIZATION

ADOPT A FLEXIBLE INFRASTRUCTURE TO SCALE AI FASTER

Different AI workloads demand different compute capabilities. Where new infrastructure is required, opting to stick with x86 architectures over ARM-based options allows enterprises to retain existing x86 applications, which simplifies AI scale-outs. Organizations can further accelerate deployment timelines by leveraging pre-optimized libraries, containers, and reference implementations that run quickly on their chosen infrastructure.

The compute profile of that infrastructure will be then defined by enterprise AI requirements. Deep learning is highly data intensive and demands the greater memory bandwidth and parallel processing of GPUs. At a small scale, most inference tasks can comfortably be handled using the higher computational efficiency and task orchestration strengths of CPUs. Combining the CPU's strengths with the GPU's intrinsic parallelization offers a pathway to handle the largest models and meet expanding AI demands.

Such flexible, scalable infrastructure not only ensures near-term performance gains but also positions organizations to respond to future AI breakthroughs, without requiring a complete technology overhaul.

AMD INSTINCT™ ACCELERATORS: HIGH PERFORMANCE AT ANY SCALE

AMD Instinct™ MI325X accelerators match competitors for 8x GPU training performance[5] while achieving up to 1.4x greater inference performance.[6] Backed by a rapidly expanding AI software stack in ROCm™, AMD Instinct™ accelerators support seamless out-of-the-box execution for over one million models.

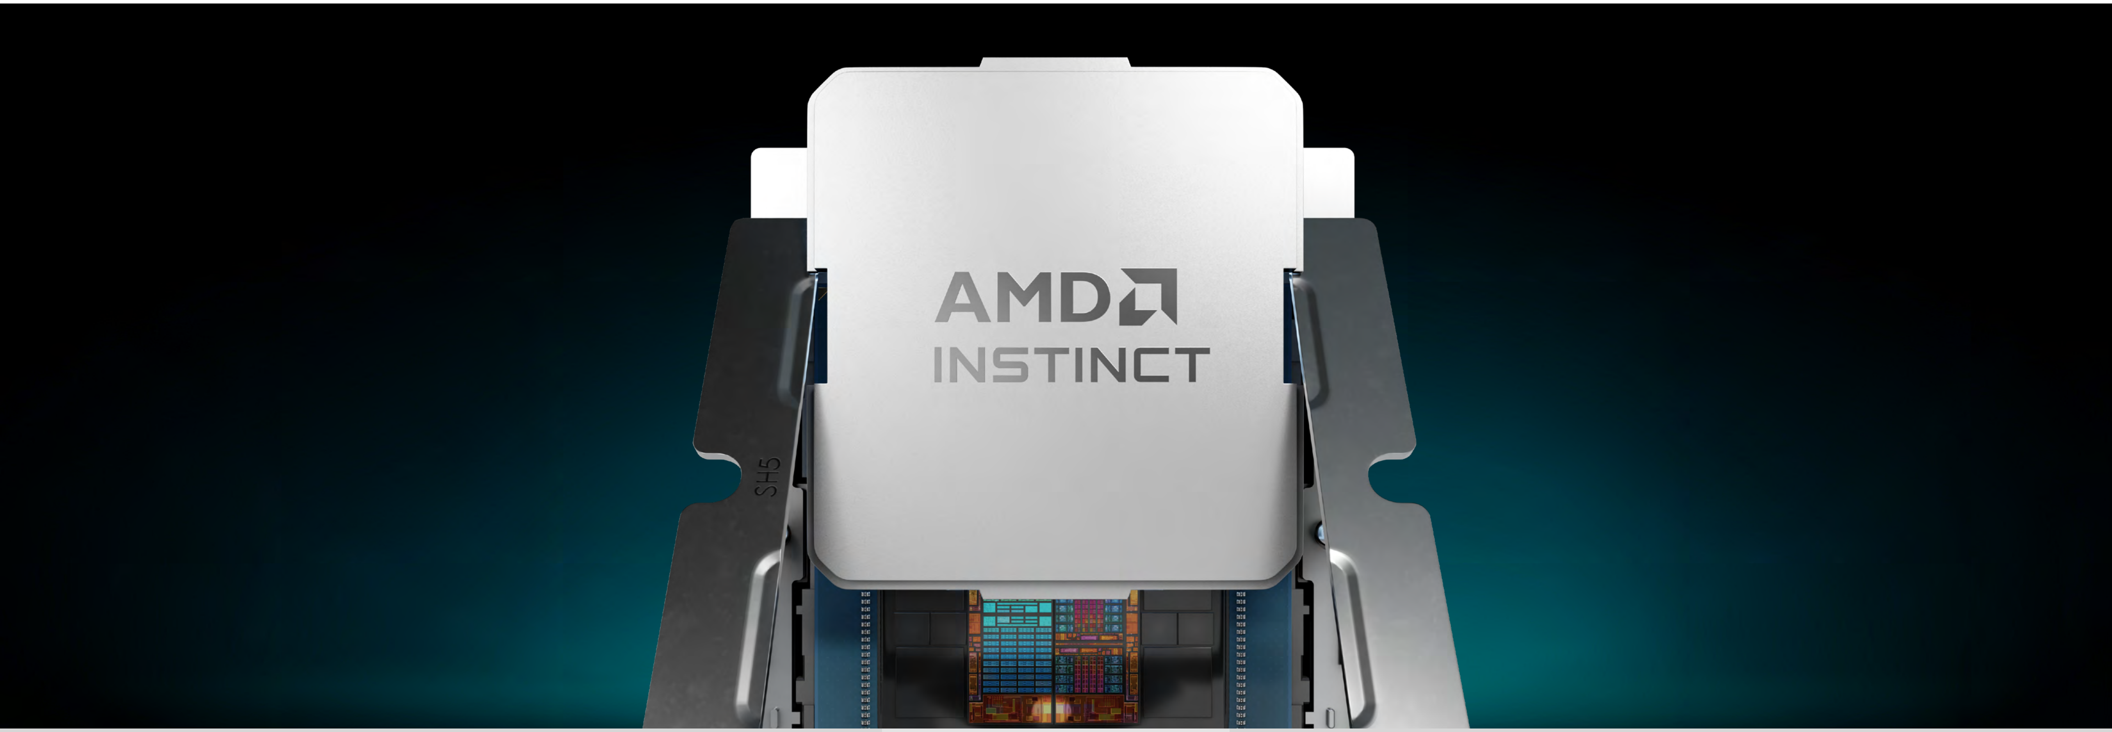

AMD

KEY INFRASTRUCTURE AREAS FOR AI MODERNIZATION

DRIVE GREATER NETWORK EFFICIENCY AND STREAMLINED MANAGEMENT

The high throughput and performance sensitivity of AI workloads places tremendous pressure on data center networks.

Network Interface Cards (NICs) and Data Processing Units (DPUs) are crucial tools for alleviating the strain AI workloads put on backend networks.

Enterprises can leverage NICs incorporating additional processing resources to offload tasks from the host CPU. This offload can include packet classification, encryption/decryption, data-plane acceleration, and more advanced congestion management. Meanwhile, DPUs featuring programmable cores can run networking, security, or storage services, further reducing CPU load. They also allow direct data movement between network, storage, and GPUs/CPUs without heavy host involvement.

AMD PENSANDO™ NETWORKING: ADVANCING PERFORMANCE IN AI NETWORKS

AMD Pensando™ Salina 400, a fully P4 programmable DPU, provides 2x the bandwidth, connections per second, packets per second, and storage operations compared to previous generations.[7] And the AMD Pensando Pollara 400 NIC uses UEC-ready RDMA, which has a 6x faster message completion time and a 5x faster collective completion time than RoCEv2.

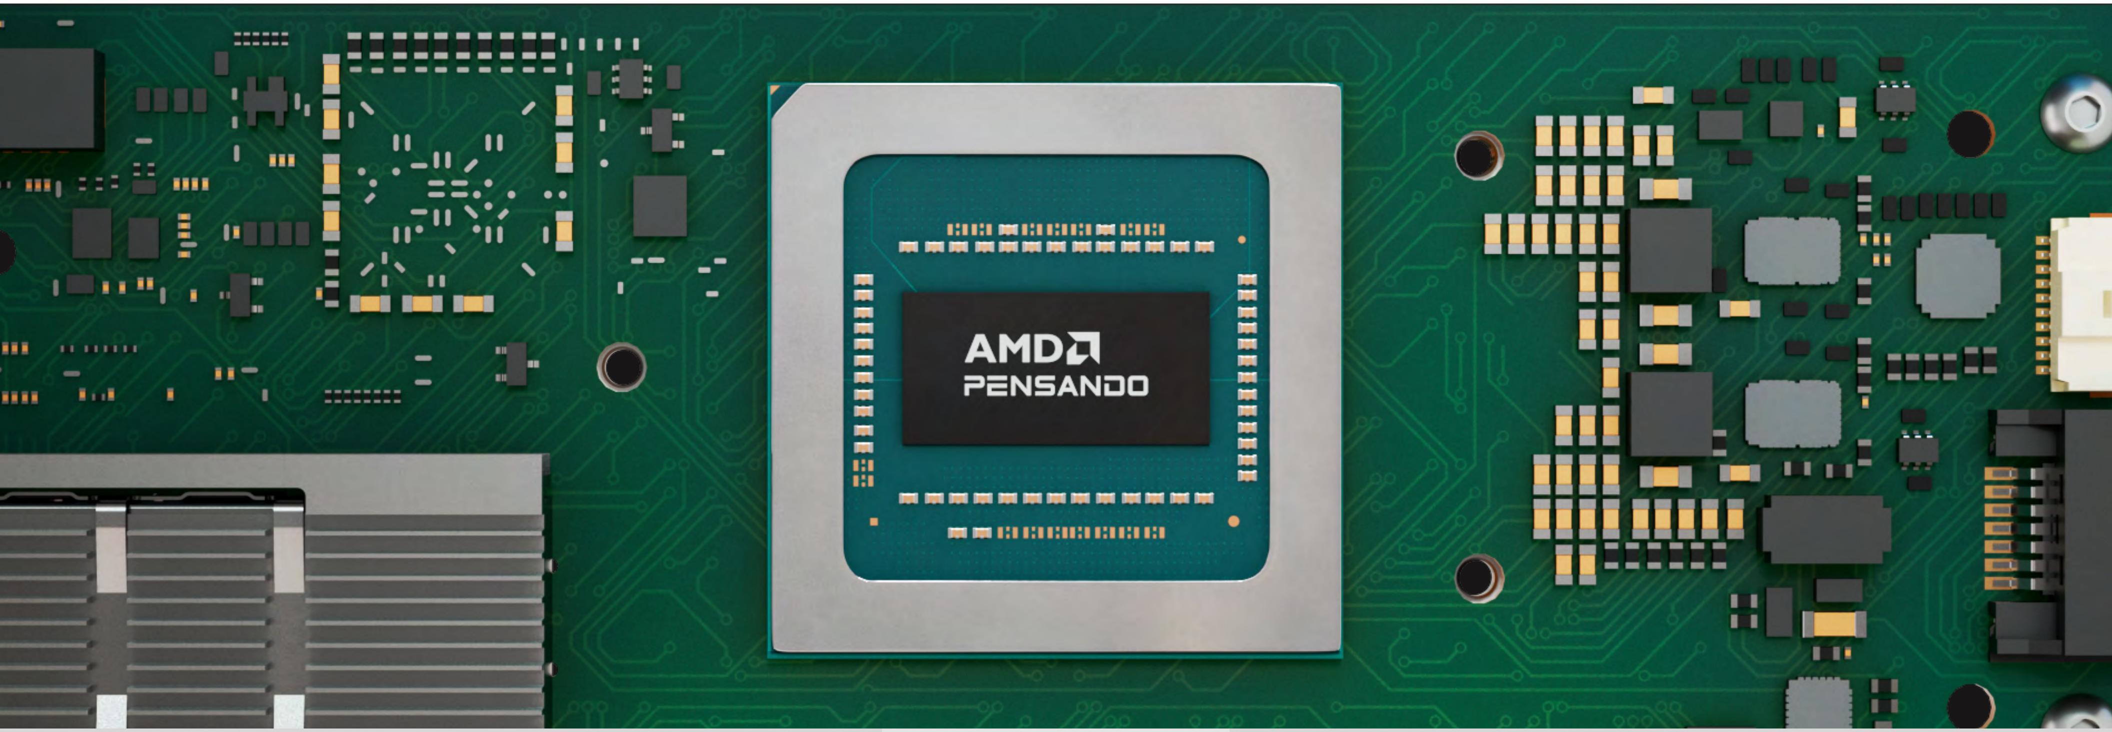

AMD

KEY INFRASTRUCTURE AREAS FOR AI MODERNIZATION

BEYOND THE DATA CENTER

AI is not reserved for large scale cloud and data center server deployments. For many end users, AI PCs will provide many day-to-day business benefits.

By processing data locally rather than relying on cloud resources, these AI PCs can deliver real-time performance, even when network connections aren't optimal. This approach also keeps personal or sensitive data on the device, enhancing privacy protections. Business users benefit from AI-assisted apps that streamline productivity tasks, automate data processing, and provide personalized insights based on usage patterns, all at lightning-fast speeds.

AMD RYZEN™ PROCESSORS: LEADING THE AI PC ERA

AMD Ryzen™ AI 300 Series processors support hundreds of different AI experiences and ship in a wide range of Microsoft Copilot+ PC laptops available today. These systems offer up to 1.4x the multithreaded performance compared to competitive offerings[8] and up to 23 hours of multiday battery life.[9]

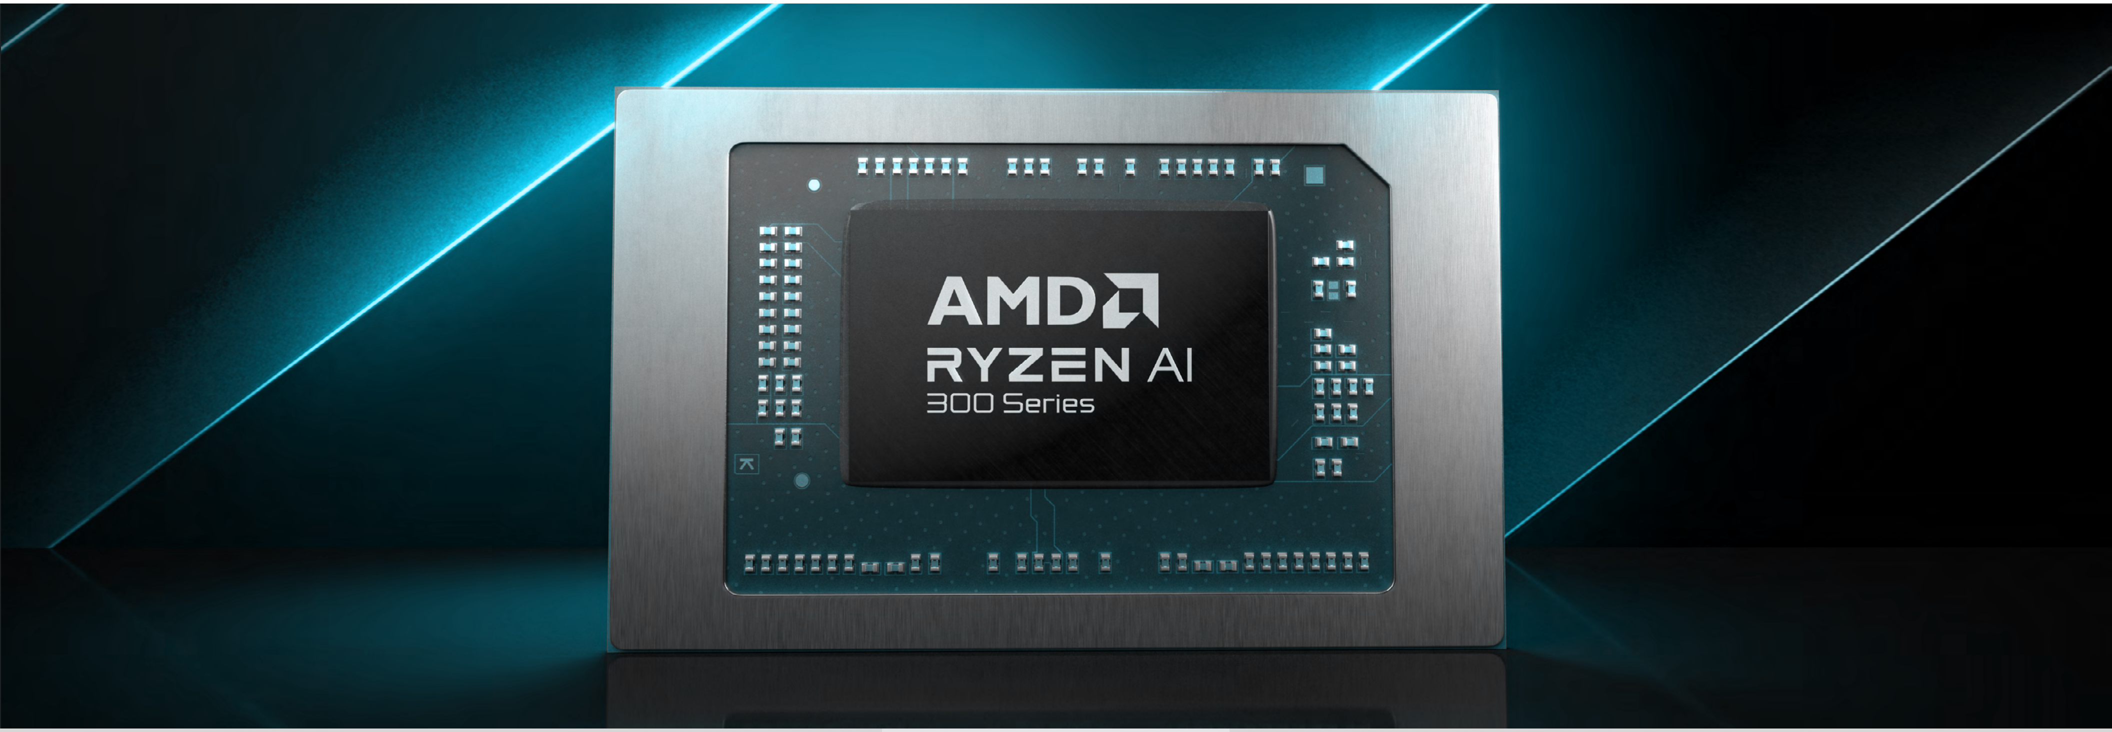

ACHIEVE SUSTAINED SUCCESS WITH AMD

AMD: THE ENTERPRISE PARTNER FOR KEEPING PACE WITH AI'S CONTINUOUS EVOLUTION

AI is being embedded across the enterprise. It will be tailored to specific tasks, varying by use case and department, and increasingly specialized by domain and industry.

The AMD end-to-end portfolio of AI solutions provides every enterprise with a pathway to scaling and accelerating time-to-results on AI adoption. It also offers tight control over ongoing operational costs, due to AMD EPYC CPUs having excellent performance per watt. That adds up to lower space and power utilization, and lower licensing costs.[10] Or, to put it another way: sustained success.

AMD is committed to maintaining an open ecosystem to help enterprise partners keep pace with the rapid evolution of the AI landscape. AMD customers are not subject to proprietary vendor lock-ins and can take advantage of AI advancements regardless of their source. This open ecosystem approach also means simplified validation and integration, day zero support from key partners, and consistent alignment with regulatory guidelines and industry best practices.

PLAN FOR AI WITH CONFIDENCE

Long-term planning is key to success for any enterprise data project. Partnering with AMD means working with a stable, decades old technology leader that continually invests in research and development, and has a long history of reliably delivering on product roadmaps and performance goals. This reliable execution gives enterprise customers peace of mind that their infrastructure investments and plans rely on a stable, dedicated technology partner.

Get in touch to discuss how working with AMD can accelerate business outcomes with AI for your organization, and deliver scalable, sustained success.

Let's talk

CLAIMS:

1. 9xx5TCO-002A: 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=SP9xxTCO-002A.
2. 9xx5-012: TPCxAI @SF30 Multi-Instance 32C Instance Size throughput results based on AMD internal testing as of 09/05/2024 running multiple VM instances. The aggregate end-to-end AI throughput test is derived from the TPCx-AI benchmark and as such is not comparable to published TPCx-AI results, as the end-to-end AI throughput test results do not comply with the TPCx-AI Specification.
   2P AMD EPYC 9965 (384 Total Cores), 12 32C instances, NPS1, 1.5TB 24x64GB DDR5-6400 (at 6000 MT/s), 1DPC, 1.0 Gbps NetXtreme BCM5720 Gigabit Ethernet PCIe, 3.5 TB Samsung MZWLO3T8HCLS-00A07 NVMe®, Ubuntu® 22.04.4 LTS, 6.8.0-40-generic (tuned-adm profile throughput-performance, ulimit -l 198096812, ulimit -n 1024, ulimit -s 8192), BIOS RVOT1000C (SMT=off, Determinism=Power, Turbo Boost=Enabled)
   2P AMD EPYC 9755 (256 Total Cores), 8 32C instances, NPS1, 1.5TB 24x64GB DDR5-6400 (at 6000 MT/s), 1DPC, 1.0 Gbps NetXtreme BCM5720 Gigabit Ethernet PCIe, 3.5 TB Samsung MZWLO3T8HCLS-00A07 NVMe®, Ubuntu 22.04.4 LTS, 6.8.0-40-generic (tuned-adm profile throughput-performance, ulimit -l 198096812, ulimit -n 1024, ulimit -s 8192), BIOS RVOT0090F (SMT=off, Determinism=Power, Turbo Boost=Enabled)
   2P AMD EPYC 9654 (192 Total cores) 6 32C instances, NPS1, 1.5TB 24x64GB DDR5-4800, 1DPC, 2 x 1.92 TB Samsung MZQL21T9HCJR-00A07 NVMe, Ubuntu 22.04.3 LTS, BIOS 1006C (SMT=off, Determinism=Power)
   Versus 2P Xeon Platinum 8592+ (128 Total Cores), 4 32C instances, AMX On, 1TB 16x64GB DDR5-5600, 1DPC, 1.0 Gbps NetXtreme BCM5719 Gigabit Ethernet PCIe, 3.84 TB KIOXIA KCMYXRUG3T84 NVMe, , Ubuntu 22.04.4 LTS, 6.5.0-35 generic (tuned-adm profile throughput-performance, ulimit -l 132065548, ulimit -n 1024, ulimit -s 8192), BIOS ESE122V (SMT=off, Determinism=Power, Turbo Boost = Enabled)
   Results:
   CPU Median Relative Generational
   Turin 192C, 12 Inst 6067.531 3.775 2.278
   Turin 128C, 8 Inst 4091.85 2.546 1.536
   Genoa 96C, 6 Inst 2663.14 1.657 1
   EMR 64C, 4 Inst 1607.417 1 NA
   Results may vary due to factors including system configurations, software versions and BIOS settings. TPC, TPC Benchmark and TPC-C are trademarks of the Transaction Processing Performance Council.
3. 9xx5TCO-001B: 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 39100 units of SPECrate2017_int_base performance as of October 10, 2024. This scenario 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 an 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'.
4. 9xx5-040A : XGBoost (Runs/Hour) throughput results based on AMD internal testing as of 09/05/2024.
   XGBoost Configurations: v2.2.1, Higgs Data Set, 32 Core Instances, FP32
   2P AMD EPYC 9965 (384 Total Cores), 12 x 32 core instances, 1.5TB 24x64GB DDR5-6400 (at 6000 MT/s), 1.0 Gbps NetXtreme BCM5720 Gigabit Ethernet PCIe, 3.5 TB Samsung MZWLO3T8HCLS-00A07 NVMe®, Ubuntu® 22.04.4 LTS, 6.8.0-45-generic (tuned-adm profile throughput-performance, ulimit -l 198078840, ulimit -n 1024, ulimit -s 8192), BIOS RVOT1000C (SMT=off, Determinism=Power, Turbo Boost=Enabled), NPS=1
   2P AMD EPYC 9755 (256 Total Cores), 1.5TB 24x64GB DDR5-6400 (at 6000 MT/s), 1DPC, 1.0 Gbps NetXtreme BCM5720 Gigabit Ethernet PCIe, 3.5 TB Samsung MZWLO3T8HCLS-00A07 NVMe®, Ubuntu 22.04.4 LTS, 6.8.0-40-generic (tuned-adm profile throughput-performance, ulimit -l 198094956, ulimit -n 1024, ulimit -s 8192), BIOS RVOT0090F (SMT=off, Determinism=Power, Turbo Boost=Enabled), NPS=1
   2P AMD EPYC 9654 (192 Total cores), 1.5TB 24x64GB DDR5-4800, 1DPC, 2 x 1.92 TB Samsung MZQL21T9HCJR-00A07 NVMe®, Ubuntu 22.04.4 LTS, 6.8.0-40-generic (tuned-adm profile throughput-performance, ulimit -l 198120988, ulimit -n 1024, ulimit -s 8192), BIOS TTI100BA (SMT=off, Determinism=Power), NPS=1
   Versus 2P Xeon Platinum 8592+ (128 Total Cores), AMX On, 1TB 16x64GB DDR5-5600, 1DPC, 1.0 Gbps NetXtreme BCM5719 Gigabit Ethernet PCIe, 3.84 TB KIOXIA KCMYXRUG3T84 NVMe®, Ubuntu 22.04.4 LTS, 6.5.0-35 generic (tuned-adm profile throughput-performance, ulimit -l 132065548, ulimit -n 1024, ulimit -s 8192), BIOS ESE122V (SMT=off, Determinism=Power, Turbo Boost = Enabled)
   Results:
   CPU Run 1 Run 2 Run 3 Median Relative Throughput Generational
   2P Turin 192C, NPS1 1565.217 1537.367 1553.957 1553.957 3 2.41
   2P Turin 128C, NPS1 1103.448 1138.34 1111.969 1111.969 2.147 1.725
   2P Genoa 96C, NPS1 662.577 644.776 640.95 644.776 1.245 1
   2P EMR 64C 517.986 421.053 553.846 517.986 1 NA
   Results may vary due to factors including system configurations, software versions and BIOS settings.

together we advance_

CLAIMS: *(continued)*

5. MI325-012: Overall GPU-normalized Training Throughput (processed tokens per second) for text generation using the Llama2-7b chat model running Megatron-LM v0.12 (BF16) when using a maximum sequence length of 4096 tokens comparison based on AMD internal testing as of 10/4/2024. Batch size according to largest micro-batch that fits in GPU memory for each system. AMD Instinct batch size 8, Nvidia batch size 2.
   Configurations:
   AMD Development system: 1P AMD Ryzen 9 7950X (16-core), 1x AMD Instinct™ MI325X (256GB, 1000W) GPU, 128 GiB memory, ROCm 6.3.0 (pre-release), Ubuntu 22.04.2 LTS with Linux kernel 5.15.0-72-generic, PyTorch 2.4.0.
   Vs.
   An Nvidia DGX H200 with 2x Intel Xeon Platinum 8468 Processors, 1x Nvidia H200 (141GB, 700W) GPUs, 2 TiB (32 DIMMs, 64 GiB/DIMM), CUDA 12.6.37-1, 560.35.03, Ubuntu 22.04.5, PyTorch 2.5.0a0+872d972e41.nv24.8.
   MI325X system median 12509.82 tokens/second/GPU
   H200 system median 11824.09 tokens/second/GPU
   Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations. MI325-012
6. MI325-004: Based on testing completed on 9/28/2024 by AMD performance lab measuring text generated throughput for Mixtral-8x7B model using FP16 datatype. Test was performed using input length of 128 tokens and an output length of 4096 tokens for the following configurations of AMD Instinct™ MI325X GPU accelerator and NVIDIA H200 SXM GPU accelerator.
   1x MI325X at 1000W with vLLM performance: 4598 (Output tokens / sec)
   Vs.
   1x H200 at 700W with TensorRT-LLM: 2700.7 (Output tokens / sec)
   Configurations:
   AMD Instinct™ MI325X reference platform:
   1x AMD Ryzen™ 9 7950X CPU, 1x AMD Instinct MI325X (256GiB, 1000W) GPU, Ubuntu® 22.04, and ROCm™ 6.3 pre-release
   Vs
   NVIDIA H200 HGX platform:
   Supermicro SuperServer with 2x Intel Xeon® Platinum 8468 Processors, 8x Nvidia H200 (140GB, 700W) GPUs [only 1 GPU was used in this test], Ubuntu 22.04) CUDA® 12.6
   Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations.
7. PEN-012: Measurements conducted by AMD Performance Labs as of Aug 27, 2024 on the current specification for the AMD Pensando™ Salina DPU accelerator designed with AMD Pensando™ 5nm process technology, projected to result in delivering 400Gb/s line-rate estimated performance.
   Estimated delivered results calculated for AMD Pensando™ Elba DPU designed with AMD Pensando 7nm process technology resulted in 200Gb/s line-rate performance.
   Actual results based on production silicon may vary.
   Salina projected performance:
   Bandwidth: 400Gbps
   Connections per second: 10M
   Packets per Second: 100MPPS
   Encryption Offloads: 400 Gbps
   Storage IOPS: 4 Million
   Actual results and specifications may vary based on production silicon.
8. STXP-12: Testing as of Sept 2024 by AMD performance labs on an HP EliteBook X G1a (14in) (40W) with AMD Ryzen AI 9 HX PRO 375 processor, Radeon™ 890M graphics, 32GB of RAM, 512GB SSD, VBS=ON, Windows 11 Pro vs. a Dell Latitude 7450 with an Intel Core Ultra 7 165H processor (vPro enabled), Intel Arc Graphics, VBS=ON, 16GB RAM, 512GB NVMe SSD, Microsoft Windows 11 Pro in the application(s) (Best Performance Mode): Cinebench R24 nT. Laptop manufactures may vary configurations yielding different results. STXP-12.
9. STXP-32: Based on internal testing by AMD as of 9/23/24. Battery life results evaluated by operation of a nine-participant Microsoft Teams video conference on battery. Test configuration for AMD and Intel systems run from power level 90% > 45% @150nits brightness and power mode set to ""best power efficiency.""
   System config: HP EliteBook X G1a (14in) with an AMD Ryzen AI 9 HX PRO 375 processor (40W), Radeon™ 890M graphics, 32GB RAM, 512GB SSD, VBS=ON, Windows 11 Pro. System config: Apple MacBook Pro 14 with M3 Pro 12-core processor, Apple integrated graphics, 36GB RAM, 1TB SSD, MacOS 15.0. System Config: Dell Latitude 7450 with an Intel Core Ultra 7 165H processor (28W) (vPro enabled), Intel Arc Graphics, VBS=ON, 16GB RAM, 512GB NVMe SSD, Windows 11 Pro.
   Manufacturers may vary configurations yielding different results. Performance may also vary based on use of latest drivers. STXP-32.
10. 9xx5TCO-001C: 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 39100 units of SPECrate2017_int_base performance as of October 10, 2024. This scenario 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 an 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=epyc5#9xx5TCO-001B

AMD
together we advance_