4TH GEN AMD EPYC^{**} PROCESSORS FOR HEALTHCARE AND LIFE SCIENCES

AMD together we advance_data center-computing

Second Edition February 2024

4TH GEN AMD EPYC[™] PROCESSORS FOR HEALTHCARE AND LIFE SCIENCE

CONTENTS

INTRODUCTION 3
POWERING HCLS WORKLOADS WITH AMD EPYC 4
GENOMICS PERFORMANCE 5
DRUG DISCOVERY PERFORMANCE 5
HEALTHCARE APPLICATION PERFORMANCE 6
HIGH PERFORMANCE COMPUTING (HPC) SYSTEMS 8
ENHANCING SECURITY WITH AMD EPYC 8
OPTIMIZING TCO WITH AMD EPYC 9
Innovative Processor Design and Manufacturing9 Leadership in Energy Efficient Computing9
CONCLUSION
FOOTNOTES

INTRODUCTION

Healthcare and life sciences (HCLS) is one of the world's most dynamic and rapidly evolving industries. It encompasses a range of interconnected fields including pharmaceuticals, biotechnology, medical devices, healthcare services, and health insurance. The common theme among these diverse fields is that each is driven by innovation that constantly pushes the boundaries of what is possible to improve quality of life on the planet.

Today, HCLS finds itself at yet another inflection point with the widespread undertaking of digital transformation. HCLS Chief Information Officers (CIOs), Chief Technology Officers (CTOs), and Chief Digital Officers (CDOs) who are leading the effort find themselves inundated with persuasive business cases for new digital tools directly from their constituent business users. The sudden availability of generative artificial intelligence (GenAI) techniques has greatly accelerated the demand for such innovation.

CIO, CTO and CDO decision-making transcends buzz-worthy trends, focusing instead on fostering a synergistic partnership with the business units they support. This collaboration is aimed at effectively addressing customer needs, whether the customer is a patient or an ecosystem partner. Importantly, these leaders embrace the responsibility of fiscal stewardship, judiciously directing technology investments with a long-term goal of achieving sustainable growth and efficiency.

There are always significant challenges facing HCLS technology leaders, but today two challenges have risen to the forefront:

- Protecting valuable data assets by actively managing cybersecurity. According to the FBI, the healthcare industry was hit hardest by ransomware attacks in 2022 compared to other critical infrastructure.¹ Furthermore, healthcare organizations across the world averaged 1,463 cyberattacks per week in 2022, up 74% compared with 2021.²
- Rising healthcare costs. Globally, rising costs threaten both the accessibility and sustainability of essential medical services, with total global spend up to US \$9 trillion in 2020 (approximately 11% of global GDP).³ Of great concern to the HCLS industry is that these trends can disproportionately impact the economically vulnerable thereby widening health disparities based on ethnicity or race. Effective cost reduction in healthcare and life sciences is not just an economic imperative but a moral one too, as it can enhance the universal availability of quality healthcare, fostering a healthier, more equitable society.

The challenge then is how can HCLS technology leaders continue to drive innovation while doing their part to keep the enterprise secure and control costs? At a minimum, they must provide infrastructure that satisfies technical performance requirements, but must also balance that within financial cost and cybersecurity constraints (Figure 1).



Figure 1: HCLS technology leaders must balance these elements for success

AMD is a trusted technology partner to many industries, known for performance leadership, outstanding TCO, and security-enhancing technology. This document highlights offerings from AMD for the healthcare and life sciences sector, focusing on its pioneering solutions tailored to meet the challenges of digital transformation. Solutions from AMD are designed with performance and security features that can drive down TCO and empower leaders to navigate and thrive in the evolving digital landscape.

POWERING HCLS WORKLOADS WITH AMD EPYC

HCLS industry computational demands are as diverse as they are complex. At the heart of these challenges is the need to process and analyze vast data sets. This necessitates system processors that are adept at handling large-scale data analysis leading to rapid inferences. Near real-time data access is crucial for healthcare providers, where immediate patient treatment decisions hinge on the swift performance capabilities of their IT systems. Moreover, the computational requirements in life sciences extend to intensive tasks like genomics analysis and molecular dynamics for drug discovery.

Furthermore, a flexible approach to cloud integration is essential, with a hybrid model that supports both on-premises applications and cloud-based services, adaptable to the unique operational needs of each organization. CPUs powering the HCLS sector must therefore be robust, swift, and versatile, capable of supporting a wide range of standards and accommodating the dynamic nature of healthcare and life sciences workloads.

4th Gen AMDeration server processors, the EPYC9004 series, are an exceptional fit for these workloads, no matter where they are running, because these processors optimize compute power and

high-speed data transfer using a revolutionary microarchitecture. Two key components represent the cornerstone for this architecture: the Core Complex (CCX) and the Core Chiplet Dies (CCD). The Core Complex (CCX) is a cluster of up to 8 CPU cores with two threads per core (for multitasking) that share L3 cache. The CCX is fabricated onto the CCD, or simply called the "Chiplet." The relationship between core, CCX and CCD is shown in Figure 2, which shows the design of the AMD EPYC[™] 9700x series of CPUs. The AMD EPYC 9100 series and 9600 series of CPUs has 32MB of L3 cache per CCX.



Figure 2: Relationship between core, core complex (CCX), and die (CCD, or Chiplet).

There are two Zen microarchitecture variants of AMD EPYC processors to select from depending on intended use:

- Zen 4: This variant is optimized for high performance in general computing workloads. These are designed for high clock speeds that support single-threaded applications such as electronic health records (EHR) or similar database driven applications. "Zen 4" processors are available with up to 12 CCDs, yielding up to 96 cores. In one example, compared to the prior generation, the "Zen 4" core delivered 45% more integer and 75% more floating point performance for 64-core processors^{4.5} allowing it to manage large genomic datasets efficiently.
- **Zen4c:** This variant is optimized for floating point operation (FLOP) density and efficiency such as for cloud applications. It is designed for multi-threaded performance including genomic sequencing, GenAl large language models, and image processing. It has the same register-transfer logic as the "Zen 4" core, but its physical layout takes less space and is designed to deliver more performance per watt. "Zen 4c" processors are available with up to 8 CCDs yielding 128 cores. In both "Zen 4" variants, the shared L3 cache in the CCX aids in reducing latency and improving data access times, which is crucial for maintaining high processing efficiency, especially in multi-threaded and multi-core scenarios. Next, we examine several HCLS workloads to better understand how AMD EPYC 9004 series processors deliver optimal performance.

GENOMICS PERFORMANCE

AMD EPYC processors excel in speeding the analysis of gene sequence variance, a critical step in powering precision medicine. The processor's floating-point compute power is complemented by high memory bandwidth, allowing it to manage large genomic datasets efficiently. Additionally, the instruction set is intricately optimized for vector and matrix operations, which are commonplace in genomic computations, offering a robust platform that enhances the speed and accuracy of genomic analysis.

In benchmark testing, servers powered by AMD EPYC completed genomic variation tasks **28% faster** than alternative currentgeneration processorsvi. The results in Figure 3. This test is based on the Broad Institute's Genome Analysis Tool Kit (GATK), which is the industry standard for identifying single nucleotide polymorphism and indels in germline DNA and RNA data.⁷ Figure 3 shows GATK benchmark testing where, lower scores indicate faster processing. Again, 4th Gen AMD EPYC[™] 9654 processors can complete GATK workloads much faster (28%) than the Intel 4th generation Xeon Platinum 8490H processors in 2P configurations.



Figure 3: GATK workload elapsed time (lower is better)

DRUG DISCOVERY PERFORMANCE

AMD EPYC 9004 processors are available with a high thread capacity of up to 256 threads. Threads are an essential feature for managing the extensive number of compute tasks, such as those inherent in drug discovery research. This research often employs molecular dynamics tools and Cryo-EM to model and simulate new medicines and require a high number of compute cores as well.

One can visualize the effect of the number of cores per server socket has on the speed of executing these computations using benchtop testing. In a series of tests comparing a control processor (Intel Xeon Platinum 8480+ with 56 cores in a 2P configuration), 2P AMD EPYC 9754 processors scored higher benchmarks across three different standard tools. In these tests, a single 2-socket server was used with standard workloads across three different molecular dynamics codes. Processors contained either two 56 core processors (control), two 96 core processors (AMD EPYC 9654), or two 128 core processors (AMD EPYC 9754). In these tests, both AMD processors were able to exhibit compelling performance gains compared to the control, with **up to 2.19 times the performance**.⁸ The scores are reported relative to the control which is assigned a value of 1 as seen in Figure 4, which shows the benchtop testing results of three different molecular dynamics tools show performance gains compared to an Intel Xeon Platinum 8480+ (56 cores). In all tests, superior performance is represented by higher scores.





Another set of benchmarks demonstrate the advantages of AMD EPYC 9004 processors used in determining molecular structures of candidate drugs. **Relion 4** (short for REgularised Llkelihood OptimisatioN 4) is a software package that employs an empirical Bayesian approach for electron cryo-electron microscopy (cryo-EM) structure determination. In third party testing, when performance is normalized to control (Intel Xeon 8490H), AMD EPYC processors returned **from 1.23 times to up to 2.12 times** better performance (see Figure 5).⁹ This benchmark tested an AMD EPYC CPU variant with AMD's 3D V-Cache™, the AMD EPYC 9684X. AMD 3D V-Cache™ delivers additional cache memory on top of each CPU die, and for the 4th generation brings 1.15TB of L3 cache. AMD uses a direct copper-to-copper hybrid bonding process that enables more than 200 times the interconnect densities of current 2D technology and more than 15 times the interconnect technologies that use solder bumps for the connection.¹⁰



Figure 5: Relion v4 benchmark testing (higher is better)

HEALTHCARE APPLICATION PERFORMANCE

Hospitals and health delivery organizations require robust and highly reliable IT infrastructure. EHRs and picture archiving and communication systems (PACS) are examples of healthcare applications with high transaction volumes which rely on intense database interactions. Testing by Epic Systems Corporation, makers of the Epic EHR, highlights 34.2M Gref/s on ODB server running 2P AMD EPYC 9654 system and the results show how AMD EPYC 9004 series processors deliver excellent performance and scalability needed to successfully manage critical patient information.¹¹

4th Generation AMD EPYC processors are ideal choice for running these kinds of high-workload demand systems. In addition to high core density, EPYC processors' high-performance interconnect helps ensure outstanding scalability and communication efficiency across the CPU. Components communicate using AMD Infinity Fabric technology, which connects cores to other components in the multi-chip "Zen" architecture including memory, PCIe[®] Gen 5 I/O channels, and security mechanisms (discussed later). These enhancements in memory and I/O design can deliver better application performance by getting data to and from the processor faster and more efficiently than ever before.

In fact, the AMD EPYC 9004 Series processors demonstrate a significant leap in performance compared to prior generations. The ODB server is the heart of the Epic product suite. Epic uses Chronicles running on the InterSystems IRIS® database to store and distribute patient record information to the rest of the systems. Epic tested servers powered by various sizes of AMD EPYC 9004 Series Processors to verify their efficiency and performance and to determine the scalability of database operations. These tests measure Global References per second (GREFs/s), and the results show how AMD EPYC 9004 Series Processors deliver the excellent performance and scalability needed to successfully manage critical patient information.

Testing by **Epic Systems Corporation**, makers of the industry leading Epic EHR highlights that even the smaller 32-core 4th Gen AMD EPYC CPUs can handle approximately **30% more users** per server compared to their 32-core 3rd Gen counterparts.¹² In fact, 4th Gen AMD EPYC 9654 CPUs can handle 34.2M Global references per second and shows an impressive performance improvement of about **2.55 times** in Global References per Second (GREFs/s) over the 3rd Gen AMD EPYC 7763 CPUs. These results are indicative of the processors' enhanced capacity to support a larger number of users and more efficient application deployment in healthcare environments.

Figure 6 shows 4th Gen AMD EPYC[™] 9004 Series Processors supporting more users per server and more GREFs/s compared to 3rd Gen AMD EPYC 7003 Series Processors.



Figure 6: Generational users/server and GREFs uplifts

In addition to EHR specific tests, in standardized testing 4th Gen AMDeration processors delivered superior performance in the opensource HammerDB testing tool. AMD EPYC 9654 servers demonstrated 2.5x the performance of older generation Intel processors in the HammerDB TPROC-H, which is based on the Transaction Processing Performance Council's TPC-H benchmark that measures decision support system workloads. In the HammerDB TPROC-C test, which is based on the TPC-C benchmark for online transaction (OLTP) throughput, the same processor demonstrated **2.7x the performance** of prior generation Intel processors. In virtualized workloads, these same processors delivered **2.8x the performance** versus the competition's prior generation for VMmark[™] which measures performance, scalability, and power consumption of virtualized platforms. Figure 7 summarizes the test results of HammerDB TPROC-H¹³, TPROC-C,¹⁴ and VMware® VMmark® performance testing versus Intel Xeon processors. Higher scores indicate greater performance for all tests.



Figure 7: HammerDB TPROC-C and TPROC-H, and VMMark uplifts

HIGH PERFORMANCE COMPUTING (HPC) SYSTEMS

Some challenges in HCLS require a supercomputer to tackle. To enable exascale-class scientific research, AMD in collaboration with the U.S. Department of Energy, Oak Ridge National Laboratory, and Cray Inc., has created the Frontier HPC system. Frontier is a next generation supercomputer, delivering more than 1.6 exaflops of peak processing power, all within only 74 cabinets of servers.¹⁵

Each node in Frontier is built using one optimized 3rd generation AMD EPYC processor and four purpose built AMD Instinct[™] 250X graphical processing unit (GPU), all interconnected using AMD Infinity Fabric. Biological research teams use Frontier while investigating protein engineering to enable new types of industry, design new pharmaceuticals and medical therapies. AMD EPYC[™] processors and AMD Instinct[™] accelerators continue to be the solutions of choice for many of the most innovative, energy efficient and fastest supercomputers in the world.

AMD now powers 140 supercomputers on the <u>November 2023</u>. <u>Top500 list</u>, representing a 39% year-over-year increase, as well as eight of the top ten most energy efficient supercomputers in the world based on the November 2023 <u>Green500 list</u>.¹⁶ AMD is bringing these supercomputing technologies to our datacenter CPU and GPU products.

ENHANCING SECURITY WITH AMD EPYC

Cybersecurity is paramount in HCLS given the growing number of cyberattacks and the highly valuable nature of the data. AMD processors are designed with security built-in at the silicon level and directly integrated within the Infinity Fabric. These features, collectively known as AMD Infinity Guard, provide a unique and robust set of security that help complement industry ecosystem partners at the software and system levels.⁷⁷

Given the broad adoption of private clouds in HCLS settings, virtual server and hypervisor security is of high importance. Infinity Guard provides enhanced security for exactly those environments. Table 1 provides a summary of the features and benefits, and of special note is the ability of AMD security features to encrypt data while it is being processed within the virtual CPU.

- Secure Encrypted Virtualization (SEV): Encrypts each virtual machine's memory with a unique key managed exclusively by the processor. This unique approach to encryption helps ensure that even if a hypervisor is compromised, the data within each virtual machine remains protected.
- Secure Encrypted Virtualization-Encrypted State (SEV-ES): Protects a guest operating system from attacks on its register state from a malicious hypervisor. It encrypts all CPU register contents when a virtual machine stops running, which deters the leakage of information in CPU registers to components like the hypervisor. It can even detect malicious modifications to a CPU register state.
- Secure Encrypted Virtualization-Secure Nested Pages (SEV-SNP): Adds strong memory integrity protection to help prevent malicious hypervisor-based attacks like data replay, memory re-mapping, and more in order to create an isolated execution environment. This advanced virtualization security is particularly vital in healthcare, helping ensure that each segment of the system is securely compartmentalized.

Collectively, Infinity Guard enables principles of a "zero trust" architecture, which is increasingly adopted in the healthcare sector.¹⁸ Zero trust is a security model that operates on the premise that no entity, internal or external, should be automatically trusted, and verification is required from any entity trying to access resources in the network.

The SEV and its advanced iterations (SEV-ES and SEV-SNP) inherent in these processors underpin this strategy by helping ensure that each virtual machine is isolated, and its data encrypted, thereby enforcing strict access controls, and minimizing the risk of unauthorized data access within the network.

Overall, the Infinity Guard suite's robust security features are critical in maintaining constant vigilance over data integrity and privacy. This approach is crucial in healthcare, where patient data security and system reliability are of utmost importance.

OPTIMIZING TCO WITH AMD EPYC

AMD designed the 4th generation EPYC processor for optimal total cost of ownership considering both acquisition and long-term operational costs.

INNOVATIVE PROCESSOR DESIGN AND MANUFACTURING

AMD understands that technology for creating processor cores is evolving differently from the technology used for creating the analog circuits that connect to memory and I/O devices. AMD "Zen 4" architecture decouples cores from the I/O manufacturing process, such that in the new 4th generation EPYC processor the cores are produced with 5nm technology, and the I/O die is created with 6nm processes.

This decoupling enables innovation to proceed at different paces which yields the performance benefits already described. It also provides benefits to the manufacturing process, as multi-die (5 and 6 nm) helps reduce waste in fabrication compared to monolithic processors. In monolithic manufacturing a single flaw can cause the entire processor to be rejected. With multi-die, AMD more efficiently produces and delivers cutting edge technology faster and typically at less cost compared to monolithic manufacturing.

LEADERSHIP IN ENERGY EFFICIENT COMPUTING

AMD "Zen 4c" cores are optimized for density and efficiency. The physical layout takes less space, allowing up to 128 cores on a single processor, and is designed to deliver more performance per watt. As a result, AMD can help HCLS customers gain more computational potential per server rack without increasing power consumption or cooling costs. More cores with efficient thermal management lead to better energy efficiency, which can reduce the overall electricity costs for running and cooling the servers, a significant factor in the total cost of ownership for data centers.

As the HCLS industry advances digital transformation including GenAI, 4th generation AMD EPYC processors provide a clear path for scaling up. All AMD EPYC 9004 series processors are compatible with the Socket SP5 platform, allowing one to swap out processors if more cores are eventually needed.

In addition, higher core densities can mean more powerful compute capacity in less physical space. As an example, AMD estimated the server capacity required to run 1,995 virtual machines. Typically, one would need 15 legacy 2P Intel 8380 powered servers to run that workload, but only 5 servers based on the 4th Gen AMD EPYC processor. That's the same amount of work with one-third the servers and an estimated **52% lower power, resulting in a 40% CAPEX reduction and 60% annual OPEX reduction**.¹⁹ Lastly, edge computing has become increasingly relevant in the HCLS sector, particularly for applications that require real-time data processing while connected to a cloud-based infrastructure. Edge processing of data means that essential patient data is processed closer to its source, thereby reducing latency, but it also avoids bandwidth costs associated with data transmission to and from the cloud.

For GenAl use cases, Al inferencing often happens in real-time or near real-time, especially in applications that require immediate decisions or actions, such as in clinical decision support tools and LLM-backed health chatbots. It may be more cost effective to invest in edge resources for GenAl model execution (i.e., inferencing), and focus on cloud resources for machine learning and GenAl model training. Whether on-premises or in the cloud, AMD EPYC CPUs and AMD Instinct GPUs can provide the horsepower needed.



CONCLUSION

4th generation AMD EPYC processors represent a pivotal advancement for the HCLS industry's digital transformation. Offering a blend of speed, security features, and energy efficiency, these high-performance processors are ideally suited to meet the complex processing requirements of healthcare and life sciences workloads. By providing robust computational power, enhanced cybersecurity features, and efficient cost management, AMD innovations are crucial in empowering HCLS organizations to effectively navigate the challenges of rising costs and cybersecurity threats (see Figure 8). This positions AMD as a high priority ally in the HCLS sector's journey towards a more secure, efficient, and innovative future, fostering a healthier and more equitable society.

FOOTNOTES

- Federal Bureau of Investigation. "Internet Crime Report 2022". <u>https://www.ic3.gov/Media/PDF/AnnualReport/2022_IC3Report.pdf</u>
- Check Point Research Team. "Check Point Research Reports a 38% Increase in 2022 Global Cyberattacks", January 5, 2023 <u>https://blog.checkpoint.com/2023/01/05/38-increase-in-2022-global-cyberattacks/</u>
- World Health Organization. Global Health Expenditure Database. Accessed 11/13/2023. <u>https://apps.who.int/nha/database</u>
- 4. See AMD SP5-003B: SPECrate®2017_int_base comparison based on published scores from www.spec.org as of 03/31/2023. Comparison of published 2P AMD EPYC 9534 (1250 SPECrate®2017_int_base, 560 Total TDP W, 128 Total Cores, 2.232 Perf/W, http://spec.org/cpu2017/ results/res2023q1/cpu2017-20230116-33519.html) is 1.45x the performance of published 2P AMD EPYC 7763 (861 SPECrate®2017_int_base, 560 Total TDP W, 128 Total Cores, 1.538 Perf/ W, http://spec.org/cpu2017/results/res2021q4/cpu2017-20211121-30148.html) [at 1.45x the performance/W]. SPEC®, SPEC CPU®, and SPECrate® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information.
- 5. See AMD SP5-004B: SPECrate[®]2017_fp_base comparison based on published scores from www.spec.org as of 03/31/2023. Comparison of published 2P AMD EPYC 9534 (1160 SPECrate[®]2017_fp_base, 560 Total TDP W, 128 Total Cores, \$17606 Total CPU \$, 2.071 Perf/W, <u>http:// spec.org/cpu2017/results/res2023q1/cpu2017-20230116-33521.html</u>) is 1.75x the performance of published 2P AMD EPYC 7763 (663 SPECrate[®]2017_fp_base, 560 Total TDP W, 128 Total Cores, \$15780 Total CPU \$, 1.184 Perf/W, <u>http://spec.org/cpu2017/results/ res2021q4/cpu2017-20211121-30146.html</u>) [at 1.75x the performance/ W] [at 1.57x the performance/CPU\$]. AMD 1Ku pricing and Intel ARK.intel.com specifications and pricing as of 3/31/23. SPEC[®], SPEC CPU[®], and SPECrate[®] are registered trademarks of the Standard Performance Evaluation Corporation. See <u>www.spec.org</u> for more information.
- 6. GATK 'sample' data set is downloaded from: <u>https://</u> s3.amazonaws.com/parabricks.sample/parabricks_sample.tar.gz. 20x coverage sample HG002 is downloaded from: <u>https://</u> storage.googleapis.com/brain-genomics-public/research/ sequencing/fastq/novaseq/wgs_pcr_free/20x/HG001.novaseq.pcr-<u>free.20x.R1.fastq.gz</u>. Characterization based o26658919 sequences, ~2200MB gzip file in shorter paired-end sequence NA12878 inputs.
- The Broad Institute. Broad Institute to release Genome Analysis Toolkit 4 (GATK4) as open-source resource to accelerate research. May 24, 2017. <u>https://www.broadinstitute.org/news/broad-institute-release-genome-analysis-toolkit-4-gatk4-open-source-resource-accelerate</u>
- Source: https://www.amd.com/content/dam/amd/en/documents/ epyc-business-docs/performance-briefs/amd-epyc-9754-pbmolecular-dynamics.pdf
- Derived from OpenBenchmarking.org, <u>https://openbenchmarking.org/test/pts/relion-1.0.1</u>. Testing not independently verified by AMD.
- 10. AMD, AMD 3D V-Cache[™] Technology, <u>https://www.amd.com/en/</u> technologies/3d-v-cache
- 11. <u>https://www.vmware.com/files/pdf/solutions/siemens-syngo-pacs-</u> with-horizon-view-reference-architecture.pdf.

- AMD, Powering Electronic Healthcare Software Electronic Health Records, July 2023. <u>https://www.amd.com/content/dam/amd/en/ documents/epyc-business-docs/performance-briefs/amd-epyc-9004-pb-epic-healthcare.pdf</u>
- 13. MySQL[®] 8.0.17 OLTP comparison based on AMD measured median scores on 2P EPYC 9654 compared to 2P Xeon Platinum 8380 running virtualized HammerDB TPROC-C (KVM Hypervisor Virtualization server environment with 400 WH and 64 users) as of 12/10/2022. System configurations: 2P AMD EPYC 9654 96-Core Processor, 24 x 32GB DDR5-4800, 8 x 3.2TB (Production platform), 1 x 25GBE Mellanox Technologies MT27710 Family [ConnectX-4 Lx], BIOS RTI1002E, AMD Titanite. 2P Intel(R) Xeon(R) Platinum 8380 CPU @ 2.30GHz, 16 x 32 GB DDR4-3200, 8 x 3.84TB (Kioxia KCD6XLUL3T84), 1 x 25GBE Mellanox Technologies MT27710 Family [ConnectX-4 Lx], BIOS 1.1a Supermicro SYS-620U-TNR. Both systems used Ubuntu® 22.04.1, SMT ON, 1 container per VM, 10 VMs, each 16 vCPUs, 32GB ram, 100GB disk, HammerDB Version 4.5, MySQL Version 8.0.17, Hypervisor QEMU KVM. Results: 2x AMD EPYC 9654 (~4,851,655 TPROC-C tpm/~2,087,994 NOPM) vs. 2x Xeon Platinum 8380 (~1,788,730 TPROC-C tpm/~770,179 NOPM) for ~2.71x the tpm/NOPM performance. Results may vary. SP5-071A
- MySQL® 8.0.17 DSS comparison based on AMD measured median scores on 2P 96-core EPYC 9654 compared to 2P 40-core Xeon Platinum 8380 running virtualized HammerDB 4.5 TPROC-H SF30 (KVM Hypervisor Virtualization server environment with 4 streams, 4 virtual units, calculating throughput with 4 streams x 22 queries x 3600 divided by the slowest VU completion time in seconds), MySQL 8.0.17, as of 12/13/2022. Configurations: 2x AMD EPYC 9654 (~1469.03 TPROC-H queries/hour) vs. 2x Xeon Platinum 8380 (~587.515 TPROC-H queries/hour) for ~2.5x the QPH performance. SP5-103
- 15. Oak Ridge National Laboratory, https://www.olcf.ornl.gov/frontier/
- 16. <u>https://www.amd.com/en/newsroom/press-releases/2023-11-14-</u> latest-top500-list-highlights-several-world-s-fast.html
- AMD Infinity Guard features vary by EPYC[™] Processor generations. Infinity Guard security features must be enabled by server OEMs and/ or Cloud Service Providers to operate. Check with your OEM or provider to confirm support of these features. Learn more about Infinity Guard at <u>https://www.amd.com/en/technologies/infinity-guard</u>. GD-183
- Department of Health and Human Services, HHS Cybersecurity Program. Zero Trust in Healthcare, 10/1/2020. <u>https://www.hhs.gov/ sites/default/files/zero-trust.pdf</u>
- 19. 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 score of 1995 VMs using the performance metric based on the published scores for Intel Xeon and AMD EPYC CPU based servers. This estimation reflects a 3-year time frame. This TCO model does not include any networking or storage that is external to the servers. Since PUE, if applicable, is a function of data center hardware power, any data center power associated with networking and storage is also not included.

This analysis compares a 2P AMD EPYC_9654 powered server with a VMs score of 418 per server; link and score details - 2-node, 2P 96-core EPYC 9654 powered server running VMware ESXi 8 RTM (40.19 @ 44 tiles/836 VMs, https://www.vmware.com/content/dam/ digitalmarketing/vmware/en/pdf/vmmark/2022-11-10-HPE-ProLiant-DL385-Gen11_2-Hosts.pdf; to a 2P Intel Xeon Platinum_8380 based server with a VMs score of 133 per server; link and score details - 2node, 2P 28-core Xeon Platinum 8180 running VMware ESXi v6.5.0b (8.11 @ 8 tiles/156 VMs), https://www.vmware.com/content/dam/ digitalmarketing/vmware/en/pdf/vmmark/2021-08-24-DellEMC-PowerEdge-R750.pdf.

The Intel server chassis is 2RU with a cost of \$2500. The AMD server chassis is 2RU with a cost of \$3000. Both AMD EPYC and Intel based servers use the same cost for the following elements of the analysis: 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 \$0.16; power drop per rack of 8kW; and a PUE (power usage effectiveness) of 1.7 - PUE is Power Usage Effectiveness.

The EPYC powered solution is estimated to take: 5 total 2P EPYC_9654 powered servers at a hardware only acquisition cost of \$54782 per server, which includes \$11805 per CPU, total system memory of 3072GB, which is 16GB of memory / core and a total system memory cost using DDR5 DIMMs of \$27792; internal storage cost of \$380. The total estimated AMD EPYC hardware acquisition cost for this solution is \$273910. Each server draws ~955.57kWhr per month. For the 3-year time frame of this EPYC powered solution analysis the: total solution power cost is ~\$46784.544 which includes a PUE factor of 1.7; the total admin cost is ~\$55251, and the total real estate cost is ~\$19440, using 1 racks. A 3-year TCO estimate for the AMD solution is \$395386.

The Intel based solution is estimated to take 15 total 2P Platinum_8380 powered servers at a hardware only acquisition cost of \$30342 per server, which includes \$9359 per CPU, total system memory of 2048GB, which is 25.6GB of memory / core and a total system memory cost of \$8704; internal storage cost of \$380. The total estimated Intel hardware acquisition cost for this solution is \$455130. Each server draws ~664.3kWhr per month. For the a 3-year time frame years of this Intel based solution analysis the: total solution power cost is ~\$97572.384 which includes the PUE factor; the total admin cost is ~\$165750, and the total real estate cost is ~\$38880 using 2 racks. The total 3-year TCO estimate for the Intel solution is \$757332.384. AMD EPYC powered servers have a \$361946 lower 3-year time frame TCO. Delivering a minimum score of 1995 for VMs produces the following estimated results: the AMD EPYC solution requires 67% fewer servers [1-(AMD server count / Intel server count)]; 50% less space [1-(AMD rack count / Intel rack count)]; 52% less power [1-(AMD power cost / Intel power cost)]; providing a 48% lower 3 year TCO [1-(AMD TCO / Intel TCO)]. The AMD solution delivers a 40% lower CapEx: [1-(AMD CapEx / Intel CapEx)], AMD CapEx = \$273910 and Intel CapEx = \$455130; and a 60% lower OpEx: [1-(AMD OpEx / Intel OpEx)], AMD OpEx = \$121475.544 and Intel OpEx = \$302202.384; AMD EPYC_9654 powered servers save ~317424kWh of electricity for the a 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 ~143.86 Metric Tons of CO2 equivalents. This results in the following estimated savings based on United States data, Greenhouse Gas Emissions Avoided of one of the following: - 31 USA Passenger Cars Not Driven for 1 year; or 10 USA Passenger Cars Not Driven Annually; or 357068 Miles Driven by Avg Passenger Car; or CO2 Emissions Avoided from: 16257 Gallons of Gasoline Not Used; or 159112 Pounds of Coal Not Burned in USA; or 18 USA Homes' Electricity Use for 1 year; or 6 USA Homes' Electricity Use Annually; or Carbon Sequestered equivalent to: 2374 Tree Seedlings Grown for 10 years in USA; or 173 Acres of USA Forests in 1 year; or 57.55 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 at https://www.epa.gov/energy/greenhousegas-equivalencies-calculator

AMD CPU pricing based on 1KU price as of Nov 2022. Intel® Xeon® Scalable CPU data and pricing from https://ark.intel.com as of Nov 2022. Memory pricing sourced online from https://memory.net/store/ on 09/28/2022. All pricing is in USD. 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. VMmark is a registered trademark of VMware in the US or other countries. Results generated by: AMD EPYC[™] Bare Metal Server & Greenhouse Gas Emission TCO Estimation Tool - version 6.40 Harini Malik is a Senior Director of Business Development at AMD. Tony Nunes Senior Product Manager at AMD

DISCLAIMERS

*Links to third party sites are provided for convenience and unless explicitly stated, AMD is not responsible for the contents of such linked sites and no endorsement is implied.

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. GD-18

©2024 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD logo, EPYC, and combinations thereof are trademarks of Advanced Micro Devices. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

