THE PERFECT COMBINATION: ULTIMATE PERFORMANCE AND ENERGY EFFICIENCY

AMD EPYC™ 9004 SERIES PROCESSORS

What if you could have both leadership performance and energy efficiency? When you choose 4th Gen AMD EPYC processors you can use up to 128 cores per processor to drive faster time to results, achieve quicker insights, and enjoy better business outcomes. For general purpose computing, systems using two 4th Gen 64-core EPYC 9354 CPUs deliver an impressive performance-per-watt improvement compared with those with two 3rd Gen 64-core EPYC 7763 CPUs, including 45% more integer[SP5-003B] and 75% more floating-point performance.[SP5-004B] Moving just a few servers to AMD EPYC CPUs can make a big impact. Imagine what moving a data center to AMD EPYC processor-powered servers could do for the planet. It's easy to switch because we take the mystery out of x86 CPU selection. Just choose the core count, frequency, and cache size you need—memory capacity, security features, I/O bandwidth, and more are included at no additional cost.

REFRESH AND CONSOLIDATE

- ***Cut 17 servers down to 11***

To deliver 2000 virtual machines each with with 1 core and 8 GB of memory takes an estimated 35% fewer servers when comparing 2P servers based on 96-core AMD EPYC 9654 CPUs over 60-core Intel® Xeon® Platinum 8490H CPUs [SP5TCO-049]

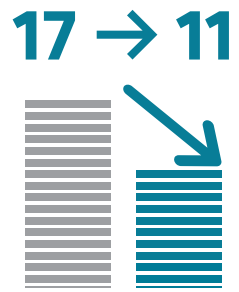

HANDLE PEAK WEB SERVER DEMANDS WITHOUT COMPROMISE

- ***Gain ~3x the cloud throughput***

Comparing two-processor servers with 128-core EPYC 9754 and those with 128-core Ampere Altra® Max M128-30CPUs on the NGINX® WRK web server benchmark[1]

ADVANCE HIGH-PERFORMANCE COMPUTING PERFORMANCE

- ***Gain 49% more floating-point performance***

Comparing servers with two 96-core EPYC 9684X processors to those with two 50-core Xeon 8490H CPUs[SP5-009E]

CUT CO_2 EQUIVALENT EMISSIONS

- ***Save up to an est. 109 metric tons over three years***

Using 11 2-socket servers with 96-core EPYC 9654 CPUs instead of 17 2-socket servers with 60-core Xeon 8490H CPUs to power 2000 virtual machines, you can reduce power by an estimated 36% annually, which is ~294 MWh less electricity used over three years[SP5TCO-049]

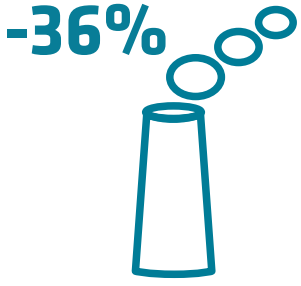

ADVANCE ENERGY-EFFICIENCY GOALS

- ***With the most energy-efficent servers***

AMD EPYC processors power the most energy-efficient servers, delivering exceptional performance, and helping reduce energy costs[SP5-028C]

ADVANCE YOUR IT INVESTMENT

- ***With 24% more integer and 25% more floating-point performance per core***

Comparing 2-socket servers with 32-core EPYC 9374F CPUs to 2-socket servers with 32-core Xeon 8462Y+ CPUs [SP5-016B, SP5-025D]

AMD EPYC 9004 SERIES ARCHITECTURE

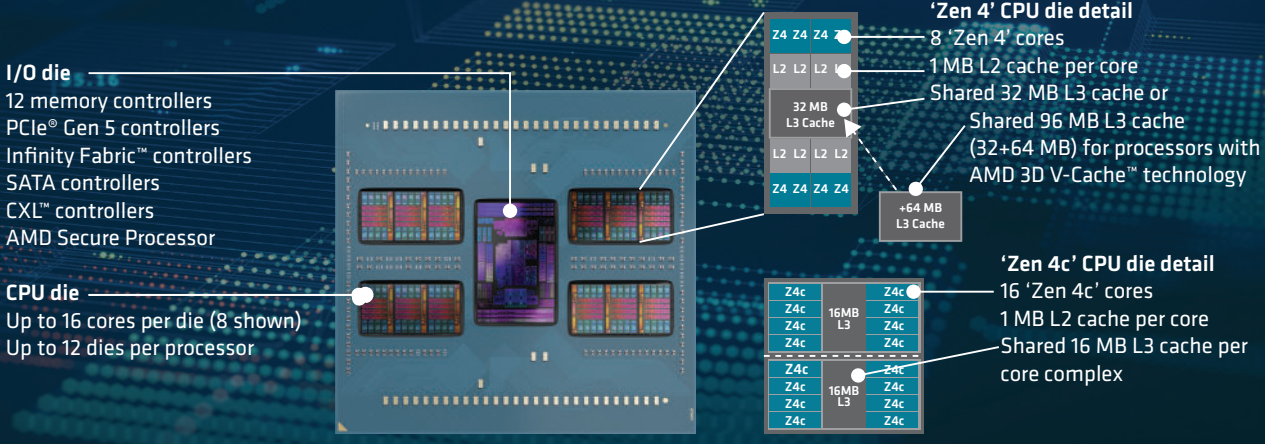

SYSTEM-ON-CHIP FEATURES

- Up to 128 PCIe Gen 5 lanes in a 1P configuration; up to 160 lanes in a 2P configuration, and up to 12 'bonus' PCIe Gen 3 lanes in a 2P configuration
- Of the 128 PCIe lanes, up to 32 are configurable as on-chip SATA controllers to access massive disk capacity
- Of the 128 PCIe lanes, up to 64 are configurable as CXL™ 1.1+ for cache-coherent memory expansion
- 12 memory controllers with capability for up to two DIMMs per channel to yield up to 6 TB of main memory per CPU—50% more than any other x86 processor[EPYC-033A]
- 2, 4, 6, 8, 10, and 12-channel memory interleaving to help optimize for both small and large memory configurations
- Integrated security processor that supports confidential computing with features including secure root of trust, secure memory encryption(SME), and secure encrypted virtualization (SEV)[2]
- 'Zen 4' core that delivers up to ~14% more instructions per clock compared to our previous generation and is further "hardened at the core" to help protect your most valuable asset—your data[EPYC-038]
- New AVX-512 instructions with 256-bit data paths support higher frequencies than our prior generation, plus they include VNNI and BFLOAT16 to match industry standards.

ADVANCING PROCESSOR ARCHITECTURE

For three processor generations, our innovative, hybrid multi-die architecture has enabled us to leap ahead of the competition in core density and performance. We have decoupled the process technologies between CPU and I/O functions so that each can advance at its own pace. This means delivering more innovations to your data center quickly. Our innovation delivers the world's highest-performing server CPU, with the 128-core EPYC 9754 CPU,[SP5-143A] and the highest x86 integer performance per core with the 16-core EPYC 9174F CPU.[SP5-014B]

AMD INFINITY ARCHITECTURE

The heart of the AMD Infinity Architecture is a leadership interconnect that supports extraordinary levels of scale at every layer. Components communicate using AMD Infinity Fabric™ technology—a connection that is used within cores, between cores, and with off-chip components—to connect 'Zen 4' processor cores, memory, PCIe® Gen 5 I/O, and security mechanisms. As a result, the architecture delivers breakthrough performance and efficiency to deliver on the promise of next-generation computing.

SYSTEM-ON-CHIP DESIGN

Our all-in philosophy means that every offering in our product line has the same built-in features. This takes the mystery out of CPU selection. Just choose the core count, cache size, and frequency your workload requires, and the rest are included at no extra cost. Our system-on-chip design means that servers don't have to include additional chip sets to handle I/O, helping reduce complexity and power consumption (see sidebar).

POWERED BY 'ZEN 4' AND 'ZEN 4C' CORES

We constantly improve our CPU design to deliver more instructions per cycle, implement new instructions, improve energy efficiency, and incorporate more hardening against security attacks. Our multi-die architecture helps us deploy different CPU designs to address workload demands without redesigning the I/O die that sits at the center of the processor. CPU dies based on 'Zen 4' cores can be used to create processors up to 96 cores. We can layer on AMD 3D V-Cache™ memory to deliver up to 1152 MB of L3 cache. And we can substitute 16-core 'Zen 4c' CPU dies to create processors with up to 128 cores.

ADVANCING EFFICIENCY

Performance and efficiency are the defining metrics of our time. In today's world, you need high performance to get the job done, but that's no longer enough. You need to do it all while using less energy to help reduce your data center's carbon footprint.

Moving just a few servers to AMD EPYC CPUs can make a big impact. For example, it takes just six 2-socket, 96-core AMD EPYC 9654 processor-powered servers to deliver 10,000 units of integer performance compared to eleven 2-socket, 60-core Xeon 8490H powered servers. The AMD EPYC solution uses an estimated 45% fewer servers and 45% less power, saving ~235,666 kWH of electricity over three years, which is ~106.81 metric tons of CO_2 equivalents.[SP5TCO-032] Imagine what moving your data center to AMD EPYC processor-powered servers can do for the planet. And more: how many servers are there in the world? What could be the benefit for the planet? Switching to EPYC CPUs could be epic.

Our 4th Gen processors power the most energy efficient servers, delivering exceptional performance and helping reduce energy costs.[EPYC-028C] For general-purpose computing, a 2-socket, 96-core EPYC 9654 CPU-based server delivers 1.8x the overall ssj_ops/W of a 2-socket, 60-core Xeon 8490H CPU-based server running the SPECpower_ssj®2008 benchmark.[SP5-011E] A 2-socket EPYC 9754 server delivers 2.7x the energy efficiency compared to a 2P Ampere Altra Max M128-30 server and 2x the energy efficiency vs. a 2P Xeon Platinum 8490H-powered server.[2]

More efficiency means more flexibility. You can use fewer servers for the same work; or accomplish more work with the same number of servers and get the job done faster. Whatever you decide, EPYC CPUs help make your data center more efficient and help you achieve your sustainability goals.

ADVANCING SECURITY

Physical and virtual threats pose a risk throughout your organization and extend to your customers. AMD Infinity Guard[GD-183] security features are a synergy between hardening in the 'Zen 4' CPU core and the dedicated, embedded security processor designed to help maintain a secure compute environment from power-on to run time.

The AMD Secure Processor scrutinizes the boot process and helps manage up to 1006 unique encryption keys known only to the security processor. Combined, these technologies help decrease potential attack surfaces as software is booted and executed and processes your data.

Our "hardened-at-the-core" security is the outcome of a continuous improvement process. For example, when we released our 2nd Gen processors with AMD Secure Memory Encryption (SME) and Secure Encrypted Virtualization (SEV), the processor architecture incorporated cache memory tagging that associates memory to the thread that 'owns' it, helping to repel side-channel attacks. Over time, we have increased the number of encryption contexts in SEV so you can accommodate more confidential virtual machines. We have helped protect virtual machine states from malicious or compromised hypervisors with AMD Secure Encrypted State (SEV-ES), and we have added AMD Secure Nested Paging (SEV-SNP) to help prevent attacks such as data replay, memory remapping, and more—all to help create confidential, isolated execution environments for virtual machines. With the larger physical memory enabled by 4th Gen AMD EPYC processors, we have increased the page table depth that can be encrypted.

ADVANCING REAL WORKLOAD PERFORMANCE

The AMD Infinity Architecture helps make real workloads run exceptionally well on servers with AMD EPYC processors. Because we view the CPU die as a modular unit of innovation, we can bake different optimizations into our processors to meet your computing needs. Whether you need to accelerate computation, speed access to data, or help defend against ever-changing security threats, there is a 4th Gen AMD EPYC processor that is perfectly suited to help drive better, energy-efficient business outcomes with modern security features that can defend against internal and external threats to keep your data safe. Our modular approach advances real workload performance in virtually every category:

- **BALANCED WORKLOADS:** We propel our mainstream product line with CPU dies made up of eight 'Zen 4' cores per die for up to 96 cores per processor. These server CPUs drive application development, business applications, data management and analytics, collaborative, and infrastructure applications. For example, a 2P, 96-core EPYC 9654 CPU-based server supports twice the number of benchmark users compared a server with two 60-core Xeon Platinum 8490H when running the SAP® SD 2-tier benchmark.SP5-056B
- **LICENSE-COST-CHALLENGED APPLICATIONS:** When you pay by the core for software licenses, you want top performance for each one. Our range of EPYC 9x74F high-frequency processors are designed to push fewer cores to higher performance per core to help you optimize costs. Use these processors for database management systems, virtualized environments, and other applications that need high per-core performance. For example, Java® enterprise middleware runs 30 percent faster using a server with two 32-core EPYC 9374F processors compared to two 32-core Xeon 8462Y+ processors as measured by max-jOPS performance on the SPECjbb® 2015-MultiJVM workload.SP5-067A
- **MEMORY-INTENSIVE WORKLOADS:** Many technical workloads process models that require large amounts of memory, putting high demands on memory throughput and cache. These include RTL simulation, computational fluid dynamics, finite element analysis, weather forecasting, and molecular dynamics. For these workloads, our EPYC 9x84X processors with AMD 3D V-Cache technology stack up to 64 additional megabytes of L3 cache on top of each CPU die for up to 1152 MB per processor. This makes the 96-core EPYC 9684X the world's highest performance x86 server CPU for technical computing.SP5-165
- **COMPUTE-INTENSIVE WORKLOADS:** For some workloads, even 96 cores per processor may not be enough. These include cloud-native applications developed with containers, virtualized environments striving for the highest number of virtual machines or containers per server, and highly parallelized workloads including financial services, life sciences, chemistry, content rendering, and delivery. When comparing eight common cloud-native workloads, a server with two 128-core EPYC 9754 CPUs delivers up to ~2.6x (~2x average) the performance of a server with two 60-core Xeon 8490H-CPUs and up to ~3.7x (~2.8x average) the performance of a server with two 128-core Ampere Altra Max M128-3 processors.[1]

"[With AMD EPYC servers], our power consumption reduced by 50 percent. But we had ten times the capacity to grow."

CHOON BOON TAN
MANAGING DIRECTOR AND HEAD OF CLOUD ENGINEERING & SERVICES
DBS SINGAPORE[3]

YOUR DATA—ANYWHERE

As an IT practitioner, you know how important it is to achieve the best possible balance between performance and efficiency. With the revolutionary AMD Infinity Architecture that delivers efficiency, performance, memory, and security features, we can help you guard your most important assets, accelerate your workloads, and modernize your data center so that you can move at the speed of your business.

AMD EPYC processor-powered servers are everywhere, with instances available from all of the major cloud providers so that you can extend your on-premises infrastructure into the cloud with the performance and efficiency that you enjoy in the data center and at the edge. You can count on the major providers rapidly adopting this newest and best generation of AMD EPYC processors ever.

READY TO MAKE THE SWITCH?

Visit amd.com/epyc

FOOTNOTES1

For details on the footnotes used in this document, visit amd.com/en/claims/epyc and https://www.amd.com/en/claims/epyc4. For a complete list of world records see amd.com/worldrecords.

1 See https://www.amd.com/system/files/documents/amd-epyc-9754-pb-cloud-native-workloads.pdf.
2 See https://www.amd.com/system/files/documents/amd-epyc-9754-pb-spec-power.pdf.
3 See https://www.amd.com/en/case-studies/dbs-bank-ltd.
GD-183 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 https://www.amd.com/en/technologies/infinity-guard.

LE-84302-01 06/23

AMD
together we advance_data center computing

READY TO CONNECT? Visit explore.amd.com/server-newsletter/sign-up