

5TH GEN AMD EPYC™ PROCESSOR ARCHITECTURE



together we advance_data center computing

*First Edition
October 2024*

5TH GEN AMD EPYC PROCESSOR ARCHITECTURE

CONTENTS

INTRODUCTION	3
Overcome Challenges	3
Optimized Performance	3
AI Leadership	3
Enterprise Trusted	4
5th Gen AMD EPYC Processors	4
HYBRID MULTI-DIE ARCHITECTURE	5
Decoupled Innovation Paths	5
5th Gen CPU Cores	5
‘Zen 5’ Core	5
‘Zen 5c’ Core	5
Modularity Enables Innovation	6
Simplified Production	7
AMD Infinity Architecture	7
I/O Die Innovation	7
5TH GEN EPYC PROCESSOR CORES	8
‘Zen 5’ Cores in EPYC 9005 Series CPUs	8
‘Zen 5’ CCD	8
‘Zen 5c’ CCD	9
AVX-512 Instruction Enhancements	9
SYSTEM-ON-CHIP DESIGN	10
Common I/O Die Features	10
EPYC 9005 Series I/O Die Features	11
AMD Infinity Fabric Technology and the I/O Die SERDES	11
NUMA Considerations	12
Reliability, Availability, and Serviceability	13
AMD INFINITY GUARD FEATURES	14
Core-Level Security Features	14
AMD Secure Processor	14
Hardware-Validated Boot	14
Transparent Secure Memory Encryption (TSME)	14
Security Features for Virtualized Environments	14
CONCLUSION	17

INTRODUCTION



Today's IT organizations must navigate the confluence of two trends. First, the demand for more enterprise applications continues unabated. Everything from traditional online transaction processing systems to highly interactive cloud-native applications are processing more data and demanding more CPU cycles to do so. Second, the transformative impact of artificial intelligence in virtually every field is requiring new, specialized infrastructure in every type of enterprise, from research to retail.

OVERCOME CHALLENGES

We believe that organizations can overcome both challenges with a coordinated approach. Modernizing with servers having highly energy-efficient CPUs can be used to consolidate infrastructure to save both space and energy. Modernization can help you get more done with fewer cores, less energy, and fewer software licenses (see [Modernize Data Center Virtualization with AMD EPYC Processors](#)). Having less data center space devoted to existing applications opens the door to servers to handle artificial intelligence and machine learning processes. We built 5th Gen AMD EPYC processors to help organizations sail smoothly through these transitions.

Systems based on 5th Gen AMD EPYC processors can support IT initiatives from data center consolidation and modernization to increasingly demanding enterprise application needs. These systems can enable expanding AI within the enterprise while supporting business imperatives to improve energy efficiency and rein in data center sprawl through high-density support for virtualization and cloud environments. Modernizing IT infrastructure is key to freeing up the space and energy to accommodate AI and other innovative business initiatives within existing data center footprints.

OPTIMIZED PERFORMANCE

In many ways, AMD EPYC processors are the world's best data center processors for demanding enterprise applications. The AMD EPYC family has established more than 400 world records on industry-standard benchmarks¹, many of which evaluate real-world application performance. The latest generation CPUs will not disappoint. We have consistently achieved double-digit gains in instruction-per-clock-cycle (IPC) performance with each new generation, and the latest 'Zen 5' core in 5th Gen AMD EPYC processors delivers significant uplifts for ML, HPC, and enterprise workloads.^{9xx5-001}

Our efficiency-optimized 'Zen 5c' core powers the CPUs with the highest core count of any x86-architecture processors, delivering the highest core density for virtualized and cloud workloads. When allocating one virtual CPU per core, the 192-core EPYC 9965 can support one-third more virtual machines than the leading available Intel® Xeon® 6E "Sierra Forest" CPU with 144 cores. Core for core, 2-socket servers with 64-core AMD EPYC 9555 CPUs achieve up to ~40% greater performance per CPU watt than 64-core Intel Xeon 8592+ processors.^{9xx5-070}

AI LEADERSHIP

A powerful foundation for AI workflows, 5th Gen AMD EPYC processors can support a broad class of AI workloads that run efficiently on CPU-only server infrastructure, without requiring GPU acceleration. Many AI workloads—many large-language models, classical image detection, fraud analysis, decision trees, and recommendations—perform effectively on traditional CPU

architectures without the need for GPU acceleration. To accelerate performance for these workloads, we have doubled most data paths in the new 'Zen 5' core, which also includes more integer arithmetic logic units (ALUs) to process data in a wider pipeline than our prior-generation processors.


Some AI workloads demand more performance than a CPU alone can provide. Very large language models, real-time model training, and other demanding generative AI workloads often require one or more advanced GPU solutions. We have optimized some of our 5th Gen processors with high core counts combined with high frequencies for use as host CPUs in GPU-accelerated systems. These processors are designed to deliver the high performance, energy efficiency, memory, and I/O scalability needed to get the most out of your investment in GPU technology.

ENTERPRISE TRUSTED

Proven performance, high efficiency, easy x86 software compatibility—all of these reasons have prompted companies, governments, and organizations around the globe to switch to AMD EPYC processors to power their most demanding computing tasks. With the leadership AMD EPYC processor portfolio available in servers from leading system vendors, qualified with leading software packages and AI frameworks, models, and tools, servers with EPYC 9005 Series processors provide a smooth path for leading compute and AI business solutions.

5TH GEN AMD EPYC PROCESSORS

This white paper describes the processor architecture that supports 5th Gen AMD EPYC processors that enable you to branch out and address a continuously widening universe of workload demands. Our hybrid, multi-chip architecture enables us to decouple innovation paths and deliver consistently innovative, high-performance products. The 'Zen 5' and 'Zen 5c' cores represent another significant advancement from the last generation, with new support for highly complex machine-learning and inferencing applications. Our system-on-chip approach helps server vendors to accelerate their designs and get innovative products into customers' hands quickly. In addition, AMD EPYC processors were the first x86 server CPUs to include an integrated, embedded security processor that is "hardened at the core" to help secure customer data whether in a central data center or distributed across locations at the network edge. Finally, this paper will review some of the design choices that enable no-compromise single-socket servers as well as some of the most powerful two-socket servers on the planet.



HYBRID MULTI-DIE ARCHITECTURE

The most industry-influential innovation in AMD EPYC processors is the hybrid multi-die architecture first introduced with second-generation EPYC processors. We anticipated the fact that increasing core density in monolithic processor designs would become more difficult over time. One of the primary issues is the fact that the process technology for CPU cores are on different innovation paths than the technology that lays down the analog circuitry to drive external pathways to memory, I/O devices, and an optional second processor. These two technologies are linked together when creating monolithic processors and can impede the swift delivery of products to market.

DECOUPLED INNOVATION PATHS

AMD EPYC processors have decoupled the innovation paths for CPU cores and I/O functions into different dies that can be developed on their own timelines and produced with process technologies appropriate for the tasks they need to accomplish. Generation over generation, we have pushed the size of CPU dies smaller as process technology allows (Figure 1). Today's 'Zen 5' cores are produced using 4nm process technology, the 'Zen 5c' core is produced at 3nm, and the I/O die remains at 6nm from the prior generation.

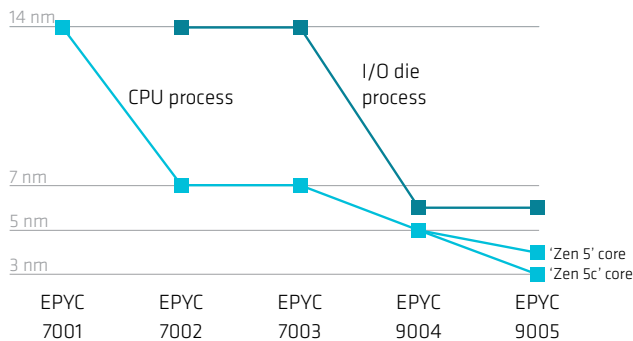


Figure 1: AMD has excelled at driving down CPU process technology by decoupling CPU and I/O innovation paths

This approach is more flexible and dynamic than trying to build all processor functions using one fabrication technology. With a modular approach, we can mix and match CPU and I/O dies to create specialized processors that closely match workload requirements. These range from high-performance processors with up to 192 cores to those for scaled-down systems needing as few as eight cores.

For CPU dies, this decoupling has enabled higher core densities over time. This, plus 'Zen' core innovations has led to double-digit gains in estimated instructions per second with every new generation. [9xx5-001](#)

5TH GEN CPU CORES

In 5th Gen AMD EPYC processors we use two different cores to address a range of workload needs by varying the type and number of cores and how we package them. These are outlined here and discussed in detail in the next chapter.

'ZEN 5' CORE

This core is optimized for high performance. Up to eight cores are combined to create a core complex (CCX) that includes a 32 MB shared L3 cache. This core complex is fabricated onto a die (CCD), up to 16 of which can be configured into an EPYC 9005 processor for up to 128 cores in the SP5 form factor. Compared to the previous generation, 5th Gen AMD EPYC processors, powered by the advanced 'Zen 5' core, along with faster memory and other key CPU improvements, provide 20% greater integer and 34% higher floating-point performance in 64-core processors operating within the same 360W TDP range. [9xx5-070](#), [9xx5-073](#)



'ZEN 5C' CORE

This core is optimized for density and efficiency. It has the same register-transfer logic as the 'Zen 5' core, but its physical layout takes less space and is designed to deliver more performance per watt. The 'Zen 5c' core complex includes up to 16 cores and a shared 32 MB L3 cache. Up to 12 of these CCDs can be combined with an I/O CCD to deliver CPUs with up to 192 cores in an SP5 form factor.


5TH GEN AMD EPYC PROCESSOR ARCHITECTURE

I/O die
 12 memory controllers
 PCIe® Gen 5 controllers
 Infinity Fabric™ controllers
 SATA controllers
 CXL™ 2.0 controllers
 AMD Secure Processor


CPU die
 Up to 16 cores per die
 Up to 16 dies per processor



AMD EPYC 9005 SERIES PROCESSORS (8-128 CORES)



AMD EPYC 9005 SERIES PROCESSORS (96-192 CORES)



MODULARITY ENABLES INNOVATION

Our modular approach enables us to deploy the CPU as a unit of innovation where we can create variants that are targeted to address specific workloads. It's a flexible unit that we can use to closely balance computing power and efficiency requirements with workloads such as the following:

- **BALANCED WORKLOADS:** We use our 'Zen 5' core to address mainstream performance needs including application development, business applications, data management and analytics, collaborative, and infrastructure applications. Up to 16 of these dies can be used to create processors with up to 128 cores.
- **LICENSE-COST-CHALLENGED WORKLOADS:** When you pay per-core software license fees, you want to get the most performance from each core. For these, and for other workloads needing high per-core performance, we have created a range of high-frequency options with fewer cores and higher clock speeds. These CPUs, with 'F' at the end of the part number, use multiple techniques to deliver higher frequencies. One is to employ higher-bin CPU dies that can run at higher frequencies. For example, the 64-core 9575F uses eight CPU dies that have been tested for a higher power budget per core, enabling high-frequency operation. The other is to employ more dies but with fewer active cores on each one in order to spread out the thermal load. This also provides more L3 cache per core with all cores active. For example, the 16-core EPYC 9175F uses 16 CPU dies, each with one core per die active. This results in 32 MB L3 cache per core.
- **ARTIFICIAL INTELLIGENCE WORKLOADS:** Whether the CPU is used for inferencing or training, our 64-core EPYC 9575F delivers a

33% higher core count than our prior generation's top frequency-optimized CPU. This high parallelism enables the opportunity for fast marshaling of data into GPUs in the role of a host node, and its high clock speed accelerates inferencing operations that do not require GPU acceleration.

- **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, weather forecasting, and molecular dynamics. Some business applications fall into this category as well, including Java® enterprise middleware. To satisfy these high memory demands, certain memory-optimized processors use multiple internal Infinity Fabric™ links between the I/O die and the CPU dies, effectively doubling the maximum theoretical memory throughput of each die. Compute Express Link (CXL®) 2.0 supports cache-coherent memory expansion and interleaving that enables large memory pools, software-managed tiered memory, and memory pooling.
- **COMPUTE-INTENSIVE WORKLOADS:** For some workloads, even 128 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 life sciences, chemistry, content rendering, and delivery. To address these needs, we combine up to 12 density-optimized 'Zen 5c' dies with 16 cores and a total of 1 MB of L2 and 32 MB of L3 cache. This brings the total core density to up to 192 cores per processor in the EPYC 9965 CPU, the highest of any x86-architecture CPU available today.

- **LOW CORE NEED APPLICATIONS:** In order to meet the full range of enterprise computing needs, we also offer 8- and 16-core processors that can satisfy the needs of applications with low compute needs. Because the product line also supports high core counts, you can use AMD EPYC processors throughout your data center.

SIMPLIFIED PRODUCTION

The hybrid, multi-die architecture can help reduce waste in the fabrication process, which helps to reduce costs. Production flaws on silicon wafers are inevitable, which is why each CPU and I/O die is tested. When many, small dies can be fabricated on a wafer, a production flaw might affect a single die that fails testing and is not integrated into a processor. If the wafer contains fewer, larger dies, a production flaw affects a larger amount of the wafer, reducing the overall yield in terms of the average number of processors produced per wafer. This can contribute to higher costs.

AMD INFINITY ARCHITECTURE

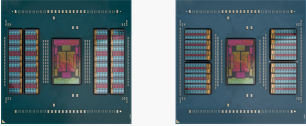
When creating a processor based on a hybrid, multi-chip architecture, the performance of the interconnect is of paramount importance. The heart of the AMD Infinity Architecture is a leadership interconnect that supports extraordinary levels of scale at every layer. Components communicate using AMD Infinity Fabric™ technology—a connection that is used between CPUs, between components in the multi-chip architecture, and to connect 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.

I/O DIE INNOVATION

While the CPU dies use our smallest process technology for high performance and low power consumption, we innovate with the I/O die to meet a different goal: interfacing the compute power of the 'Zen 5' cores to the outside world. It takes more energy to push signals to memory DIMMs, I/O devices, and second CPUs than it does to load registers and perform arithmetic calculations, so the I/O dies in 5th Gen AMD EPYC processors use slightly larger 4nm process technology.

As we have increased processing power over time, the I/O die has evolved to meet the demand. More cores require more I/O bandwidth, so the I/O die supports 12 DDR5-6000 memory controllers, PCIe Gen 5 I/O, AMD Infinity Fabric™ interconnects,

Table 1: EPYC 9005 Series At a Glance

AMD EPYC 9005 'TURIN'	
	
Core Architecture	'Zen 5' and 'Zen 5c'
CPU Process Technology	4nm and 3nm, respectively
I/O Die Process Technology	6nm
Cores	8 to 192
Performance Improvement Over Prior Generation <small>9xx5-001</small>	~37% geometric mean on ML and HPC workloads ~17% geometric mean on enterprise workloads
Max L3 Cache	512 MB
PCIe® Lanes	Up to 128 (single-socket systems) Up to 160 (2-socket servers)
Power (Configurable TDP [cTDP])	120–500W
Memory channels / max per-socket theoretical bandwidth	12 / 576 GB/s
Max Memory Capacity	6 TB DDR5-6000

SATA disk controllers, and CXL 2.0 caching accelerator and memory expansion connectivity that can be flexibly assigned to specific functions at server design time (see "CXL 2.0 Capabilities" on page 11). The I/O die is where the dedicated AMD Secure Processor resides, close to the memory controllers that manage the range of memory encryption mechanisms that are part of our AMD Infinity Guard feature set discussed in "AMD Infinity Guard Features" on page 14.

The I/O die used in 5th Gen processors has 16 Infinity Fabric connections to CPU dies. In high-core-count processors, a single connection is used to integrate each CPU die into the processor. In some processor models, two connections can be used to optimize bandwidth between the I/O die and each CPU die.

5TH GEN EPYC PROCESSOR CORES

Today's data centers need to power an ever-increasing number of applications along with a growing need to integrate AI into the business. Our focus in developing the 'Zen 5' core is to continue to accomplish double-digit percentage increases in instructions per clock cycle (IPC), and to equip the core to better handle the vast amounts of data handling and processing power that AI workloads require.

We have accomplished this goal by widening data paths in the core and using the wider paths to enable more work to be accomplished per cycle. Innovations over the 'Zen 4' core include:

- **DUAL PIPE INSTRUCTION FETCHING** along with instruction cache latency and bandwidth improvements enables parallel instruction decoding so that more instructions can be in flight at a time within the core. With highly accurate branch prediction, the processor can consume more instructions on the front end, which helps increase IPC over the 'Zen 4' core.
- **THE INTEGER PROCESSING PIPELINE** is now 8 instructions wide, one third more than the 'Zen 4' core. This increased parallelism is supported by ALUs and an improved scheduler to support a wide execution window. Integer performance is often a good predictor of business application performance, so this is an important improvement.
- **MAXIMUM DATA BANDWIDTH** has been doubled between the core and the 48 KB L1 data cache; increased data prefetch keeps data flowing into the data pipeline.
- **TO ACCELERATE AI AND HPC OPERATIONS**, the floating-point and vector ALUs have been equipped with a full, 512-bit data path. Data for AVX-512 operations can load in a single-cycle, and a total of six pipelines keep more floating-point instructions in flight at any given time.

The combined focus on highly parallel instruction decoding with increased branch prediction and a faster, wider path for data, results in a highly balanced processor architecture for accelerating AI operations along with everyday business and technical applications.

'ZEN 5' CORES IN EPYC 9005 SERIES CPUS

The 'Zen 5' core is realized with 4nm process technology through our continued deep relationships with some of the leading fabrication companies. The core is implemented with small, fast, low-power transistors for power efficiency and high performance. The 'Zen 5' core supports our hybrid, multi-chip architecture, and two versions of the core are used to meet various design goals of 5th Gen AMD EPYC processors. For example, the 'Zen 5' CPU die is optimized for high per-core performance, while the 'Zen 5c' core is optimized for high density and power efficiency. Because of our modular architecture, we can create processors that specialize in high frequency operation, giving excellent per-core performance for per-core-licensed software. We have designed the product stack so that, for a given number of cores, 5th Gen processors use approximately the same amount of power as 4th Gen CPUs, but with higher performance.

'ZEN 5' CCD

The 'Zen 5' CCD has a dedicated 1MB L2 cache per core and a shared 32MB L3 cache for up to eight cores per die (Figure 2). As a result of using smaller process technology than the 'Zen 4' CCD, more dies fit in the CPU package and can support processors with up to 128 cores. This was previously only possible with the specialized 'Zen 4c' CCD used in 4th Gen processors.





Figure 2: A 'Zen 5' CCD with 8 cores per die includes up to a 32 MB cache



'ZEN 5c' CCD

The 'Zen 5c' CPU core is designed for high density and energy efficiency. The same logic as the 'Zen 5' core is more closely packed in its core complex to build processors with up to 192 cores. The 'Zen 4c' CPU die holds 16 cores each having 1 MB L2 cache and a shared 32 MB L3 cache (Figure 3). To create processors with more than 128 cores, up to 12 of these dies can be attached to the I/O die for a total of up to 192 cores per processor for ultra-dense, high-performance systems. The 'Zen 5c' die is also used in the 128-core EPYC 9745 and the 96-core EPYC 9645 processors to provide an alternate, lower-power option compared to CPUs with the same core count built with the 'Zen 5' core die.




Figure 3: The 'Zen 5c' CPU die holds total of 16 cores per die

AVX-512 INSTRUCTION ENHANCEMENTS

Many applications today strive to gain knowledge from data, and they repeat arithmetic calculations on large amounts of data. These workloads include:

- Computational fluid dynamics
- Cryptography and data compression
- Finite element analysis
- Financial services
- Image and audio/video processing
- Life sciences

- Machine learning and inferencing
- Molecular dynamics modeling
- Oil & Gas exploration

While most applications use a single instruction to operate on a single data element (single instruction, single data [SISD]), these applications need parallel execution of multiple data elements directed by a single instruction (single instruction, multiple data [SIMD]). Some codes, including HPC, financial services, and video processing use vectors of full-precision floating-point data. Machine learning and inferencing workloads are increasingly using half-precision arithmetic including 16-bit floating point and 8-bit integer operations to speed the flow of data and reduce the power needed to process large data sets.

AVX-512 is a set of instructions based on an SIMD model. As its name suggests, a single instruction operates on a 512-bit vector of 8-, 16-, 32-, or 64-bit data values. 5th GEN EPYC processors implement the full set of AVX-512 instructions used in 4th Gen Intel Xeon processors except for FP16 data types. Our implementation of these data-heavy instructions supports the most current AVX-512 instructions and EVEX prefixes, which means code written for early AVX-512 implementations will run with no modifications.

We have expanded data paths to 512-bits starting with the 'Zen 5' core. Along with this, we have increased the floating-point queue, schedulers, and pipes. By adding more integer ALUs, the data can be read into the CPU in a single clock cycle, and processed through the pipeline expediently. If power efficiency is a priority over performance, BIOS settings can direct the processor to execute two 256-bit vectors in sequential clock cycles for AVX-512 instructions.

These expanded data paths contribute to dramatically higher performance across processor generations as measured by the HPL benchmark. Comparing 4th and 5th Gen AMD EPYC processors, two-socket servers with 192-core 5th Gen EPYC 9965 CPUs outperform 96-core 4th Gen EPYC 9654 CPUs by ~2.75x. [9xx5-080](#)

SYSTEM-ON-CHIP DESIGN

AMD EPYC processors are systems on chip (SOCs) by nature of the fact that they provide all the features needed to design complete servers without the use of a chipset. This approach can help reduce server design complexity and power consumption because fewer chips outside of the CPU are needed. These features are implemented by the I/O die (Figure 4) that ties together the CPU dies and handles all interfaces with the outside world such as memory access and I/O.

Our all-in philosophy means that every processor series offers the same built-in features. This approach takes the mystery out of CPU selection. Just choose the series, core count, frequency, and L3 cache size your workload requires, and all of the features for that series are included at no extra cost.

COMMON I/O DIE FEATURES

Just as we share the same ‘Zen 5’ logic between the ‘Zen 5’ and the ‘Zen 5c’ CPU dies, we share logic between I/O die implementations for consistency between features.




Figure 4: The I/O die implements many functions that would otherwise require external chip sets

- DDR5 MEMORY CONTROLLERS**—Twelve in the EPYC 9005 Series, Having higher performance x86 CPU cores, and more of them, creates a higher demand for memory, and more memory channels drive high memory bandwidth that keeps this equation in balance. New in the EPYC 9005 Series is support for DDR5-6000 speeds, helping as well to increase memory throughput. Memory interleaving on 2, 4, 6, 8, 10, and 12 channels helps optimize for both small- and large-memory configurations. The memory controllers include inline encryption engines for implementing AMD Infinity Guard features discussed in the next chapter.
- SMART DATA CACHE INJECTION**— This feature moves data directly from a PCIe endpoint into the L2 cache, avoiding an initial write into main memory and then a read by the CPU into its cache. This helps boost performance for networking, real-time processing, low-latency financial-services applications. This feature maintains memory and cache coherency by knowing that the data block must eventually be written to main memory.
- PCIe GEN 5 I/O**—Single-socket server configurations have up to 128 PCIe Gen 5 lanes. PCIe I/O is implemented with serializer/deserializer logic that can assume different functions, including Infinity Fabric connectivity, on-chip SATA disk controllers, and CXL 2.0 connectivity. These I/O features are subject to the constraints described in the series-specific sections that follow.
- INTERNAL INFINITY FABRIC INTERFACES** connect the I/O die with each CPU die using a total of 16 36 Gb/s Infinity Fabric links. (This is known internally as the Global Memory Interface [GMI] and is labeled this way in many figures.) In memory-speed-optimized EPYC 9005 Series processors two links connect to each CPU die for up to 72 Gb/s of connectivity.
- INTEGRATED AMD SECURE PROCESSOR** supports features including secure root of trust, transparent secure memory encryption (TSME), and, secure encrypted virtualization (SEV). See “AMD Infinity Guard Features” on page 14 for more details.
- A SERVER CONTROLLER HUB** can minimize the required chip set for basic server control functions. Depending on the I/O die, the SCH can include direct USB connectivity, 1 Gb/s LAN-on-motherboard, and various UART and I₂C and I₃C bus connectivity.

- **A SYSTEM MANAGEMENT UNIT** controls power distribution to the I/O die, CPU dies, and cores, maintaining the processor within its thermal design parameters or those set through power-management settings such as configurable thermal design power (cTDP).

EPYC 9005 SERIES I/O DIE FEATURES

The EPYC 9005 Series I/O die includes the common I/O die features plus the following additional capabilities:

- **AMD INFINITY FABRIC** is used for interprocessor communication in 2P configurations. Rather than invent new connectivity mechanisms that can delay time to market, we use the same physical interfaces for Infinity Fabric connections as for PCIe Gen 5 I/O, with different protocols layered on the physical (PHY) layer. This affords server designers the freedom to trade off more PCIe I/O lanes in exchange for fewer interprocessor communication links. We support the use of three or four links, each of which correspond to 32 Gb/s x16 PCIe speeds. When three links between processors are used, an additional 16 PCIe lanes on each CPU are available for general I/O, bringing the total I/O capacity up to 160 lanes. When four links are configured, they can support a maximum theoretical bandwidth of 512 GB/s between processors.
- **FLEXIBLE PCIe GEN 5 LANES**— Up to 32 lanes can be configured as on-chip SATA controllers, and up to 64 lanes can be configured to support CXL 2.0 connectivity. The 5th Gen Trusted I/O feature (see “AMD Infinity Guard Features” on page 14) leverages industry-standard PCIe link-encryption technology to extend trust boundaries in the cloud to establish trust in a device and its configuration.
- **UP TO 12 PCIe GEN 3 ‘BONUS’ LANES** in a 2-socket EPYC 9005 Series configuration, or 8 lanes in a single-socket configuration. In server designs, the bonus lanes are often used for access to performance-insensitive I/O such as to M.2 drives used for system boot.
- **CXL 2.0 CAPABILITIES**— The EPYC 9005 Series supports Type 1 and Type 3 device use cases including persistent memory, software-managed tiered memory, and memory pooling. The CPU can interleave memory between CXL devices and also between CXL and main memory. CXL Type 1, 2, and 3 devices are supported depending on ecosystem readiness; type 2 devices supported for proof of concept only.
- **MEMORY QUALITY-OF-SERVICE TELEMETRY**— this can provide information on the relative performance of DRAM compared to CXL memory so that operating systems and hypervisors can make informed choices on memory allocation and placement.

AMD INFINITY FABRIC TECHNOLOGY AND THE I/O DIE SERDES

A single physical layer supports I/O functions including PCIe Gen 5, AMD Infinity Fabric, SATA disk controllers, and CXL connectivity. This approach reflects our philosophy of using industry-standard, well-understood technologies that offer server designers flexibility to design innovative systems, and it simplifies our CPU designs compared to inventing proprietary interconnects.

I/O serializer-deserializer (SERDES) logic resides on the I/O die with one set of traces programmable to support multiple functions. The I/O die in EPYC 9005 provides eight 16-lane SERDES devices that provide up to 128 lanes of I/O connectivity subject to the constraints listed in Table 2.

The lanes in each SERDES can be bifurcated given constraints described in each processor’s server design documentation. Each SERDES has specific constraints, for example some are restricted to PCIe and Infinity Fabric connectivity, while others enable a richer set of functions.

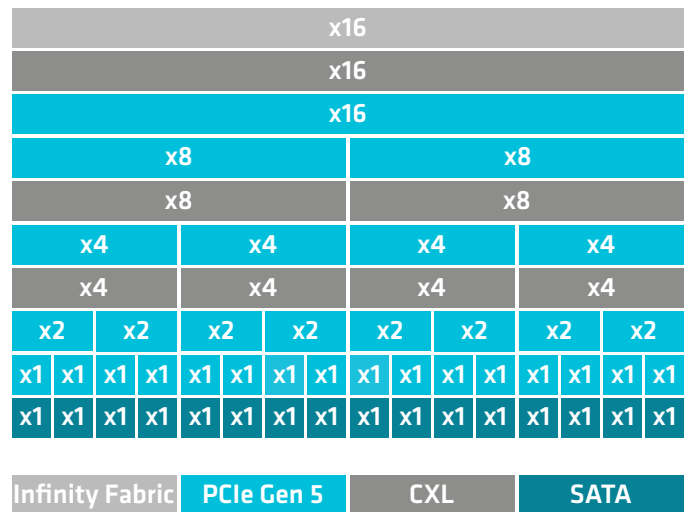


Figure 5: Idealized example of SERDES lane bifurcation options

An idealized bifurcation diagram—no single SERDES provides all of them—is illustrated in Figure 5, indicating that the entire 16-lane port can be dedicated to 16 lanes of Infinity Fabric, PCIe, or CXL connectivity. PCIe I/O can be broken down to various combinations of x8, x4, x2, and x1 bandwidth. CXL connections can have x16, x8, or x4 bandwidth. SATA controllers are x1, however if they share connectivity with PCIe on the same SERDES, a maximum of eight SATA controllers can be allocated.

Table 2: SERDES support for multiple I/O functions

CPU Series	PCIe Gen 5		Infinity Fabric	SATA	CXL 2.0
	1P Config	2P Config			
Link width	x1 (min)	x1 (min)	x16	x1	x4 (min)
EPYC 9005	128 lanes + 8 bonus PCIe Gen 3	160 lanes + 12 bonus PCIe Gen 3	3-4 links	32 drives	4 x16 links

SINGLE-SOCKET SERVER CONFIGURATIONS

AMD EPYC processors without a ‘P’ suffix can be used in single-socket and 2-socket configurations. Processor part numbers with a ‘P’ suffix are optimized for single-socket servers by dedicating all SERDES links for PCIe I/O connections only. Figure 6 illustrates a single-socket configuration using an EPYC 9005 Series processor with two DIMMs per memory channel.




Figure 6: EPYC 9005 Series processor in a single-socket server configuration with both ‘P’ and ‘G’ links dedicated to PCIe connectivity

TWO-SOCKET SERVER CONFIGURATIONS

In a multiprocessor server design, flexibility of the SERDES enables the Infinity Fabric interconnects to use the same physical infrastructure of chip’s PCIe I/O. In Figure 7, these are labeled as ‘G’ and ‘P’ links, each of which support 16 lanes of PCIe Gen 5 connectivity.

In these configurations, three or four 16-lane ‘G’ links are used to connect to the second processor. For I/O-intensive server designs, three links can be used as Infinity Fabric interconnects and one additional link from each CPU can be dedicated to PCIe Gen 5 I/O, bringing the server I/O capacity to 160 lanes.

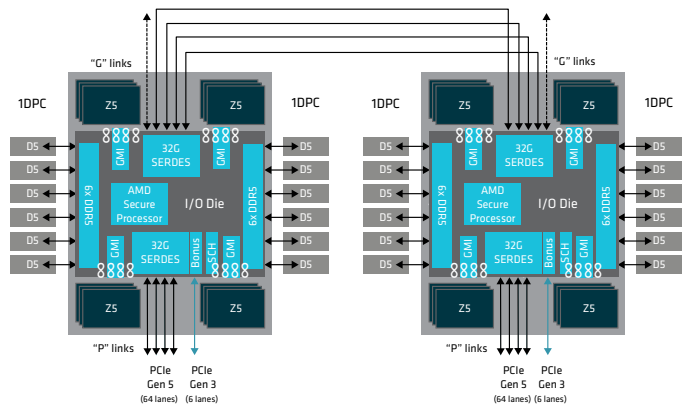


Figure 7: EPYC 9005 Series processors in a 2-socket configurations

NUMA CONSIDERATIONS

In a multi-chip architecture, there can be varying amounts of memory latency depending on the connectivity between memory controllers and CPU dies. This is known as non-uniform memory access, or NUMA. For applications needing to extract every last percent of latency out of memory accesses, they can take advantage of these varying latencies to create an affinity between specific address ranges and the CPU cores closest to that memory.

Figure 8 illustrates how this works in the EPYC 9005 Series, which is topologically equivalent to the EPYC 9004 Series. If you divide the I/O die into four quadrants for an ‘NPS=4’ configuration, you will see that six DIMMs feed into three memory controllers, which are closely connected via Infinity Fabric (GMI) to a set of up to four ‘Zen 5’, or up to three ‘Zen 5c’ CPU dies.

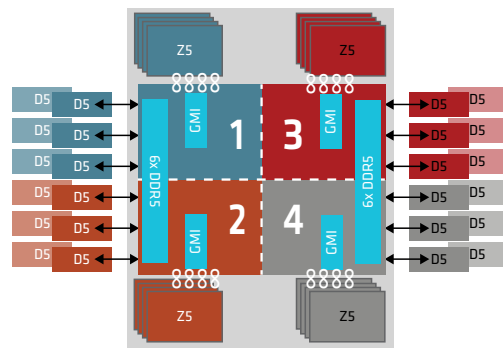


Figure 8: Dividing an EPYC 9005 Series processor into four NUMA domains can give small performance improvements for some applications

Historically, AMD EPYC 7001 Series processors located memory controllers on the same die with up to eight CPU cores, creating a tight affinity between the memory controlled by the die and the CPU cores on the die. When a memory controller had to request data destined for a different set of cores, the data had to pass from one die to another over an internal Infinity Fabric connection.

Beginning with AMD EPYC 7002 Series processors, non-uniform latency was reduced dramatically by locating memory controllers onto the I/O die. NUMA domains were flattened more by the move to 32 MB of L3 cache in the EPYC 7003 Series.

In 4th and 5th Gen EPYC processors, optimizations to the Infinity Fabric interconnects reduced latency differences even further. Today, most applications don't need to be concerned about using NUMA domains, and using the AMD EPYC processor as a single domain (NPS=1) gives excellent performance. The [AMD EPYC 9004 Architecture Overview](#) provides more details on NUMA configurations and tuning suggestions for specific applications.

RELIABILITY, AVAILABILITY, AND SERVICEABILITY

AMD EPYC processors are built with RAS features that increase their own reliability, availability, and serviceability and that of the platform in which they run.

RELIABILITY is the rate at which the CPU successfully completes work—if its calculations produce incorrect results, it becomes less than 100% reliable. AMD EPYC processors are designed with thousands of self-checking circuits that enable the processor to continue accomplishing work without interruption.

System memory is a significant source of errors in a typical platform. Two features that contribute to platform reliability are the processor's ability to overcome memory errors:

- **AMD ADVANCED MEMORY DEVICE CORRECTION (AMDC)** goes beyond standard ECC DRAM by using a type of ECC that allows large groups of bits to be corrected with negligible performance impact. Similar to Chipkill, this helps prevent a bad DRAM device on a DIMM from causing application problems. AMDC helps increase server reliability and availability by enabling DIMMs with errors to remain in service.
- **DRAM READ UECC RETRY AND DRAM ADDRESS AND COMMAND PARITY WITH REPLAY** help to overcome transient memory bus errors by replaying requests that presented errors, helping to maintain high levels of service.

AVAILABILITY is the ability of a system to accept new work, often measured as platform uptime. At the platform level, AMD EPYC processors participate in making it possible to hot swap PCIe and NVMe devices without bringing the processor down, or throttling CPU frequencies to manage DRAM thermal characteristics.

AMD EPYC processors implement the following features that contribute to platform uptime:

- **DATA POISONING** passes information on uncorrectable errors to the CPU by routing a 'poison' bit along with the data so the CPU can report the problem to the operating system with a clear indication of which process took the uncorrectable error. This prevents corrupt data from being used by applications and minimizes the scope of the error to only the process(es) depending on the data.
- **SEAMLESS FIRMWARE SERVICING (SFS) UPTIME** is a way for AMD to distribute firmware updates to large-scale operators so that they can be applied without rebooting in order to sustain the health and security of platforms. This, in turn, helps infrastructure stay up and running while still implementing firmware updates that provide for better end-user experiences. Contrast this to traditional approaches to firmware updates that involve reloading code into a SPIROM or other flash device then rebooting the platform to cause that new code to execute.

SERVICEABILITY is the capability to diagnose and repair system components. For example, AMD has observed that memory errors have become more common over time. Our advanced approach to error correction allows memory DIMMs that might otherwise cause a server outage to remain in service, helping to maintain continuous availability.

Some of the notable serviceability features in 5th Gen AMD EPYC processors include:

- **DYNAMIC POST-PACKAGE REPAIR (PPR)** enables the processor to apply a repair to a faulty DIMM without a reboot, further enabling DIMMs with errors to remain in service. This is supported on both x4 and x8 DIMMs. These repairs can be requested either in band via x86 firmware or out of band via the BMC, without any x86 firmware involvement. If a memory error occurs, the operating system is notified so that it can map out the page; dynamic PPR allows the system to avoid loss of that memory page.
- **CRASH DUMP AND MCA OVER APML** allow the processor to report both fatal and runtime errors out-of-band to the platform without the overhead of halting the x86 cores. This allows the baseboard management controller to be informed so that the platform can decide on any additional service actions.

AMD INFINITY GUARD FEATURES

Data is every organization's most precious asset, and AMD Infinity Guard hardware-based security features are designed to help protect your data from malicious users, hypervisors, and even rogue administrators. This approach can help mitigate the risks of attacks against physical DIMMs or attacks against guests in virtualized and hyperconverged environments.

Collectively, our enhancements are known as [AMD Infinity Guard security features](#). Some features vary by EPYC processor generations and/or series and 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.

CORE-LEVEL SECURITY FEATURES

Our approach to security begins in the CPU core and extends out to the system on chip and then to the platform in which the processor resides. Each 'Zen' core generation builds upon the security features of the previous one, and they mitigate vulnerabilities known at design time with no modifications necessary to application software.

The original 'Zen' core has resisted certain side-channel attacks in part because of the tagging of memory to threads when read into the processor caches. This helps reduce the possibility of one thread being able to view another thread's data when in use in the processor. In the 'Zen 4' core we introduced the option for guest operating systems in virtualized environments to run exclusively on one core—thus introducing further solutions that can help protect against side-channel attacks targeted at cached memory.

AMD SECURE PROCESSOR

Security features are managed by the AMD Secure Processor, which coordinates activities between the CPU cores, the memory controllers, and the boot firmware. The processor is a 32-bit microcontroller that runs a hardened operating system. The hardening process removes unnecessary components and applies previous security patches in the microcontroller to help reduce attack surfaces. It provides cryptographic functionality for key generation and key management, and it supervises hardware-validated boot, where the foundation for platform security starts.

HARDWARE-VALIDATED BOOT

Hardware-validated boot helps verify that the operating system or hypervisor software that you intended to load is what is actually loaded. The AMD Secure Processor loads the on-chip boot ROM that loads and authenticates the off-chip boot loader. The boot loader, in turn, authenticates the BIOS before any of the 'Zen' cores can execute the code. Once the BIOS is authenticated, the OS boot loader loads the operating system or hypervisor.

In virtualized environments, attestation helps assure that your software is loaded onto the server or service of your choice and that it is encrypted. After the boot process is complete, virtual machines asking the hypervisor for attestation can receive a checksum of the entire boot image along with a key identifying the server, location, or cloud provider. This provides evidence to guests that their software is booted without corruption on the server they intended, with encryption enabled. In cloud environments, this can help ensure that a guest virtual machine is executing where it is supposed to, in compliance with data sovereignty regulations.

TRANSPARENT SECURE MEMORY ENCRYPTION (TSME)

This foundational feature, part of AMD EPYC processors since the beginning, can be used to encrypt all of main memory with no changes required to the operating system or application software. TSME helps protect against attacks on the integrity of main memory (such as certain cold-boot attacks) because it encrypts the data. 256-bit AES-XTS encryption engines are built into the memory controllers to help reduce performance impact during reading and writing of encrypted memory. These engines can be used to encrypt memory with either 128- or 256-bit keys. The 256-bit encryption option introduced with the 'Zen 4' generation is integrated into the I/O die in order to support the United States Federal Information Processing Standards (FIPS) 140-3 standard. All of this is done without the encryption key being visible outside of the AMD Secure Processor.

SECURITY FEATURES FOR VIRTUALIZED ENVIRONMENTS

Virtualization is essential to today's IT organizations, whether they run virtualized environments in house, in the cloud, or both. Consider

the concerns of each constituency. As an IT organization or a cloud service provider, you want to ensure that a malicious virtual machine cannot penetrate the walls separating one guest from another and from the hypervisor. If a guest is able to compromise hypervisor security, there is no longer assurance that its data is safe. As a customer of a virtualized or cloud environment, you need to know not only that other guests cannot interfere with your operations and data, but also that even a malicious hypervisor (not under your control) can compromise the security of your data. Helping to secure virtualized environments is critical for cloud computing.

Across each generation of AMD EPYC processors, we have invested in increased security for virtualized environments in the form of progressively more sophisticated features that enhance the isolation between virtual machines and the hypervisor (Figure 9).

- **AMD SECURE ENCRYPTED VIRTUALIZATION (SEV)** enables hypervisors and guest virtual machines to be cryptographically isolated from one another. Thus, if malicious software is successful in evading the isolation provided by the hypervisor, or if the hypervisor itself is compromised, reading memory from another virtual machine will expose only encrypted data for which the key is stored inside of the AMD Secure Processor and memory controllers. In EPYC 9005 Series processors, up to 1006 keys can be used for virtual machine encryption. Firmware for implementing SEV is [open source](#) and subject to community inspection.
- **AMD SECURE ENCRYPTED STATE (SEV-ES)**, available beginning with 3rd Gen AMD EPYC processors, encrypts virtual machine state

when interrupts cause it to be stored in the hypervisor. With this information encrypted with the virtual machine's encryption key, a compromised hypervisor is still unable to view a virtual machine's registers.

- **AMD SECURE NESTED PAGING (SEV-SNP)**, introduced in 3rd Gen AMD EPYC processors, builds on SEV and SEV-ES by adding strong encryption to virtual machine nested page tables to help prevent attacks such as data replay, memory remapping, and more—all with the goal to create confidential, isolated execution environments for virtual machines. With the 57-bit physical memory supported beginning with the AMD EPYC 8004 and 9004 Series, we increased the page table depth to five levels.

Also included with SEV-SNP is the capability for hypervisors to use third-party encryption keys in providing attestations to guest virtual machines. Attestation contributes to confidential computing by providing users with information they need to make an informed decision regarding whether to trust the virtualized environment into which they have booted. An attestation may include measurements of the trusted computing base, including firmware, microcode, and guest characteristics including memory checksums, launch metadata, guest configuration, and x86 runtime configuration.


- **AMD SECURE MULTI-KEY ENCRYPTION (SMKE)**, introduced in EPYC 9004 Series processors, extends encryption to storage-class memory, which helps data stored on CXL-attached memory to remain encrypted across a system reboot, helping protect even persistent memory from prying eyes.

	EPYC 7001	EPYC 7002	EPYC 7003	EPYC 9004/8004	EPYC 9005
Trusted I/O					✓
SEV - Secure Nested Pages			✓	✓	✓
SEV - Encrypted State			✓	✓	✓
Secure Encrypted Virtualization (SEV)		✓	✓	✓	✓
Secure Memory Encryption (SME)	✓	✓	✓	✓	✓


Figure 9: Each new AMD EPYC processor generation delivers more features to help isolate virtual machines

-
- **AMD TRUSTED I/O**, introduced in EPYC 9005 Series processors, establishes a framework for securing virtualized I/O paths. The feature supports the PCI-SIG TDISP standard that mutually authenticates the VM and the device, then establishing an encrypted connection. Consider, for example, running ML training in a cloud service provider environment. You don't want your training data to be transferred to a GPU you don't trust, and SEV provides the framework for ensuring mutual trust and security of data in flight. Full details of trusted I/O and how it uses attestation is available in the AMD white paper [AMD SEV-TIO: Trusted I/O for Secure Encrypted Virtualization](#).

This powerful set of security features, is enabled in turn by a multi-layered set of technologies accessible by all of the major hypervisor vendors. It is an innovative set of modern security features that help decrease potential attack surfaces as software is booted, executed, and processes your data. Built-in at the silicon level, AMD Infinity Guard features offer state-of-the-art capabilities to help defend against internal and external threats. Whether yours is a small- or medium-size business or an enterprise organization, implementing robust security features on premises or in the cloud is streamlined with AMD Infinity Guard.



CONCLUSION



AMD EPYC 9005 Series processors deliver out-of-the-box performance and high core density for the growing demands of AI-enabled, business-critical data center workloads. Based on our newest 'Zen 5' core, we continue to deliver real innovation and double-digit IPC increases with each generation. Decoupling our core and I/O silicon development enables us to shrink the CPU die to 3nm in the current processor, helping improve performance along with energy efficiency.

This processor generation is targeted to deliver the balance you need to tackle both traditional enterprise and AI-enabled workloads. For traditional, virtualized, and cloud environments, up to 192 cores enable high performance through software that supports high parallelism. For AI inferencing and also AI model training, our doubling of CPU data paths since the last generation helps satisfy the voracious appetite for data that characterize these applications. Once inside the 512-bit path into the processor, an increased number of arithmetic and logic units is ready to speed the processor pipeline to deliver results quickly.

To further help protect your data in shared virtualized or cloud environments, our newest Infinity Guard feature, Trusted I/O, establishes a framework for securing virtualized I/O paths between the CPU and supported devices such as network interfaces, disk controllers, and GPU accelerators. This expands the sphere of security that protects your virtual machines to further the goal of confidential computing on premises and in the cloud.

When you deploy the EPYC 9005 series in your data center servers, you can modernize your data center and help reduce the number of servers you need for existing traditional workloads. This can free up available power and space to support increased general compute needs, or new AI initiatives, or a combination of both..

END NOTES

For details on the claims used in this document, visit amd.com/en/legal/claims/epyc.

1. AMD EPYC™ Family of Processors as of 5/02/2024. See amd.com/worldrecords for the full list
- 9xx5-001: Based on AMD internal testing as of 9/10/2024, geomean performance improvement (IPC) at fixed-frequency.
- 5th Gen EPYC CPU Enterprise and Cloud Server Workloads generational IPC Uplift of 1.170x (geomean) using a select set of 36 workloads and is the geomean of estimated scores for total and all subsets of SPECrate®2017_int_base (geomean), estimated scores for total and all subsets of SPECrate®2017_fp_base (geomean), scores for Server Side Java multi instance max ops/sec, representative Cloud Server workloads (geomean), and representative Enterprise server workloads (geomean).
 - "Genoa" Config (all NPS1): EPYC 9654 BIOS TQZ1005D 12c12t (1c1t/CCD in 12+1), FF 3GHz, 12x DDR5-4800 (2Rx4 64GB), 32Gbps xGMI;
 - "Turin" config (all NPS1): EPYC 9V45 BIOS RVOT1000F 12c12t (1c1t/CCD in 12+1), FF 3GHz, 12x DDR5-6000 (2Rx4 64GB), 32Gbps xGMI
 - Utilizing Performance Determinism and the Performance governor on Ubuntu® 22.04 w/ 6.8.0-40-generic kernel OS for all workloads.
 - 5th Gen EPYC generational ML/HPC Server Workloads IPC Uplift of 1.369x (geomean) using a select set of 24 workloads and is the geomean of representative ML Server Workloads (geomean), and representative HPC Server Workloads (geomean).
 - "Genoa Config (all NPS1): "Genoa" config: EPYC 9654 BIOS TQZ1005D 12c12t (1c1t/CCD in 12+1), FF 3GHz, 12x DDR5-4800 (2Rx4 64GB), 32Gbps xGMI;
 - "Turin" config (all NPS1): EPYC 9V45 BIOS RVOT1000F 12c12t (1c1t/CCD in 12+1), FF 3GHz, 12x DDR5-6000 (2Rx4 64GB), 32Gbps xGMI
 - Utilizing Performance Determinism and the Performance governor on Ubuntu 22.04 w/ 6.8.0-40-generic kernel OS for all workloads except LAMMPS, HPCG, NAMD, OpenFOAM, Gromacs which utilize 24.04 w/ 6.8.0-40-generic kernel.
- 9xx5-070: SPECrate®2017_int_base comparison based on published scores from www.spec.org as of 10/10/2024.
- 2P AMD EPYC 9555 (1610 SPECrate®2017_int_base, 128 Total Cores, 360W TDP, \$9,826.00 CPU \$), 4.472 SPECrate®2017_int_base/CPU W, 0.164 SPECrate®2017_int_base/CPU \$, <https://www.spec.org/cpu2017/results/res2024q4/cpu2017-20240920-44764.html>)
 - 2P AMD EPYC 9554 (1340 SPECrate®2017_int_base, 128 Total Cores, 360W TDP, \$9,087.00 CPU \$), 3.722 SPECrate®2017_int_base/CPU W, 0.147 SPECrate®2017_int_base/CPU \$, <https://www.spec.org/cpu2017/results/res2023q2/cpu2017-20230327-34992.html>)
 - 2P Intel Xeon Platinum 8592+ (1120 SPECrate®2017_int_base, 128 Total Cores, 350W TDP, \$11,600 CPU \$) 3.229 SPECrate®2017_int_base/CPU W, 0.097 SPECrate®2017_int_base/CPU \$, <https://www.spec.org/cpu2017/results/res2023q4/cpu2017-20231127-40064.html>)
- 9xx5-073: SPECrate®2017_fp_base comparison based on published scores from www.spec.org as of 10/10/2024.
- 2P AMD EPYC 9555 (1670 SPECrate®2017_fp_base, 128 Total Cores, 360W TDP, \$9,826.00 CPU \$), 4.639 SPECrate®2017_fp_base CPU W, 0.170 SPECrate®2017_fp_base/CPU \$, <https://www.spec.org/cpu2017/results/res2024q4/cpu2017-20240920-44764.html>)
 - 2P AMD EPYC 9554 (1250 SPECrate®2017_fp_base, 128 Total Cores, 360W TDP, \$9,087.00 CPU \$), 3.472 SPECrate®2017_fp_base/CPU W, 0.138 SPECrate®2017_fp_base/CPU \$, <https://www.spec.org/cpu2017/results/res2024q1/cpu2017-20240129-40783.html>)
 - 2P Intel Xeon Platinum 8592+ (1260 SPECrate®2017_fp_base, 128 Total Cores, 350W TDP, \$11,600 CPU \$) 3.600 SPECrate®2017_fp_base/CPU W, 0.109 SPECrate®2017_fp_base/CPU \$, <https://www.spec.org/cpu2017/results/res2024q3/cpu2017-20240701-43949.html>)
- 9xx5-080: AMD testing as of 09/18/2024. The detailed results show the average uplift of the performance metric (TFLOPS) of this benchmark for a 2P 192-Core AMD EPYC™ 9965 powered system compared to a 2P 96-Core AMD EPYC™ 9654 powered system running select tests on Open-Source HPL v2.3. Uplifts for the performance metric normalized to the 96-Core AMD EPYC™ 9654 follow for each benchmark:
- * hpl: ~2.75x
- CPU: 2P 192-Core AMD EPYC™ 9965 (384 total cores)
Memory: 24x 64 GB DDR5-6000
Storage: SAMSUNG MZWLO3T8HCLS-00A07
Platform and BIOS: VOLCANO RVOT1000C
BIOS Options:
SMT=Off
NPS=4
Power Determinism Mode
OS: rhel 9.4 5.14.0-42716.1.el9_4.x86_64
Kernel Options:
amd_iommu=on iommu=pt mitigations=off
Runtime Options:
cpupower idle-set -d 2
cpupower frequency-set -g performance
echo 3 > /proc/sys/vm/drop_caches
echo 0 > /proc/sys/kernel/nmi_watchdog
echo 0 > /proc/sys/kernel/numa_balancing
echo 0 > /proc/sys/kernel/randomize_va_space
echo 'always' > /sys/kernel/mm/transparent_hugepage/enabled
echo 'always' > /sys/kernel/mm/transparent_hugepage/defrag
- CPU: 2P 96-Core AMD EPYC™ 9654 (192 total cores)
Memory: 24x 64 GB DDR5-4800
Storage: SAMSUNG MZQL21T9HCJR-00A07
Platform and BIOS: Titanite_4G RT1009C

```
BIOS Options:
SMT=Off
NPS=4
Power Determinism Mode
OS: rhel 9.4 5.14.0-42716.1.el9_4.x86_64
Kernel Options:
amd_iommu=on iommu=pt mitigations=off
Runtime Options:
cpupower idle-set -d 2
cpupower frequency-set -g performance
echo 3 > /proc/sys/vm/drop_caches
echo 0 > /proc/sys/kernel/nmi_watchdog
echo 0 > /proc/sys/kernel/numa_balancing
echo 0 > /proc/sys/kernel/randomize_va_space
echo 'always' > /sys/kernel/mm/transparent_hugepage/enabled
echo 'always' > /sys/kernel/mm/transparent_hugepage/defrag
```

Results may vary based on factors including but not limited to system configurations, software versions, and BIOS settings.

© 2022–2024 Advanced Micro Devices, Inc. All rights reserved. All rights reserved. AMD, the AMD Arrow logo, EPYC, Infinity Fabric, and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. CXL is a trademark of Compute Express Link Consortium, Inc. Intel and Xeon are trademarks of Intel Corporation or its subsidiaries. SPEC® and SPECrate® are registered trademarks for Standard Performance Evaluation Corporation. Learn more at [spec.org](https://www.spec.org). Other names are for informational purposes only and may be trademarks of their respective owners. Certain AMD technologies may require third-party enablement or activation. Supported features may vary by operating system. Please confirm with the system manufacturer for specific features. No technology or product can be completely secure
LE-91401-00 10/24



Contact your Connection Account Team for more information.

Business Solutions
1.800.800.0014

Enterprise Solutions
1.800.369.1047

Public Sector Solutions
1.800.800.0019

www.connection.com/AMD