



AMD EPYC™ 9005 PROCESSOR ARCHITECTURE OVERVIEW

WHITEPAPER | 2025

TABLE OF CONTENTS

Chapter 1 - Introduction.....	3
General Specifications	3
Operating Systems	3
Chapter 2 - AMD EPYC™ 9005 Series Processors	4
“Zen 5” Core	4
Core Complex (CCX)	4
Core Complex Die (CCD)	5
I/O Die (Infinity Fabric™)	5
Memory, I/O, and Conectivity	6
Chapter 3 - NUMA Topology	7
NUMA Settings.....	7
Dual-Socket Configurations.....	9
Chapter 4 - Processor Identification	10
CPUID Instruction	10
Chapter 5 - Additional Information	10
AVX-512	10
Chapter 6 - Resources	11

CHAPTER 1: INTRODUCTION

AMD EPYC™ 9005 Series Processors represent the fifth generation of AMD EPYC server-class processors. This generation of AMD EPYC processors feature AMD's latest "Zen 5" based compute cores, next-generation I/O Die, enhanced security features, and increased memory & I/O bandwidth and speeds, all while using the existing SP5 socket/packaging.

GENERAL SPECIFICATIONS

AMD EPYC 9005 Series Processors offer a variety of configurations with varying numbers of cores, Thermal Design Points (TDPs), frequencies, cache sizes, etc. that complement AMD's existing server portfolio with further improvements to performance, power efficiency, and value for a variety of environments and workloads.. Table 1-1 highlights some features of AMD EPYC 9005 Series Processors.

PROCESSOR AT A GLANCE		
CODENAME	Turin	
COMPUTE CORES	Zen 5-based	
SOCKET	SP5	
FAMILY	Family 1Ah	
MODEL	00-0Fh	10h-1Fh
ZEN 5 CORE	"Zen5" classic	"Zen5c" dense
MAX NUMBER OF CORES (SMT THREADS)	128 (256)	192 (384)
MAX CORES PER CCX	8	16
MAX SHARED L3 CACHE SIZE PER CCX	32 MB	
NUMBER OF CORE COMPLEXES (CCXS) PER CCD	1	
MAX NUMBER OF CORE COMPLEX DIES (CCDS)	16	12
MAX CCDS PER QUADRANT	4	3
MAX # OF MEMORY CHANNELS	12 DDR5	
MAX MEMORY SPEED	6000 MT/s DDR5	
MAX # OF DIMMS PER CHANNEL (DPC)	2	
MAX MEMORY CAPACITY	8TB, based on 6TB DDR (256GB DIMMs * 2 DPC * 12 channels) + 2TB CXL (256GB CXL card * 8 cards)	
MAX LANES PERIPHERAL COMPONENT INTERCONNECT	128 lanes PCIe® Gen 5 (+ 8 lanes PCIe Gen 3)	
MAX LANES COMPUTE EXPRESS LINKS	64 lanes CXL 2.0 (4 x 16 P-links)	

Table 1-1: Common features of all AMD EPYC 9005 Series Processors

OPERATING SYSTEMS

AMD recommends using the latest available targeted OS version and updates. Please see **AMD EPYC™ Processors Minimum Operating System (OS) Versions** for detailed OS version information.

CHAPTER 2: AMD EPYC™ 9005 SERIES PROCESSORS

AMD EPYC 9005 Series Processors incorporate compute cores, memory controllers, I/O controllers, RAS (Reliability, Availability, and Serviceability), and security features into an integrated System on a Chip (SoC). The AMD EPYC 9005 Series Processor retains the proven Multi-Chip Module (MCM) Chiplet architecture of prior successful AMD EPYC processors while making further improvements to the SoC components.

The SoC includes the Core Complex Dies (CCDs), which contain Core Complexes (CCXs), which contain the “Zen 5”-based cores. The CCDs surround the central high-speed I/O Die (and interconnect via the Infinity Fabric). The following sections describe each of these components.




Figure 2-1: AMD EPYC 9005 processor with central I/O Die and 16 Zen 5 Classic-based Core Complex Dies




Figure 2-2: AMD EPYC 9005 processor with central I/O Die and 12 Zen 5c “Dense”-based Core Complex Dies

CORE COMPLEX (CCX)

The term “Core Complex” refers to a set of cores sharing a last-level (L3) cache. An AMD EPYC 9005 CCX will either have up to eight Zen5 (classic) or sixteen Zen5c (dense) cores sharing a 32MB L3 cache. Enabling SMT allows a single CCX to support twice that number, i.e., 16 or 32 concurrent hardware threads.




Figure 2-3: CCX with eight Zen5 compute cores and 32 MB shared L3 cache (corresponds to Figure 2-1, above)




Figure 2-4: CCX with sixteen Zen5c compute cores and 32 MB shared L3 cache (corresponds to Figure 2-2, above)

CORE COMPLEX DIE (CCD)

Each Core Complex Die (CCD) of an AMD EPYC 9xx5 Series Processor contains a single CCX. Cores can potentially be disabled in BIOS using one or both of the following approaches:

- Reduce the cores per CCD while keeping the number of CCDs constant. This approach increases the effective L3 cache per core ratio but reduces the number of cores sharing the L3 cache.
- Reduce the number of active CCDs while keeping the cores per CCD constant. This approach maintains the advantages of cache sharing between the cores while maintaining the same cache per core ratio.

I/O DIE (INFINITY FABRIC™)

The CCDs connect to memory, I/O, and each other's CCxs through the I/O Die (IOD). This central AMD Infinity Fabric™ provides the data path and control to interconnect CCxs, memory, and I/O. Each CCD connects to the IOD via a dedicated high-speed Global Memory Interconnect (GMI) link. The IOD helps maintain cache coherency and additionally provides the interface to extend the data fabric to a potential second processor via its xGMI, or G-links. AMD EPYC 9005 Series Processors support up to 4 xGMI (G-links) with speeds up to 32Gbps. The IOD also exposes DDR5 memory interfaces in addition to the Global Memory Interconnect (GMI) links, as well as PCIe® Gen5 and CXL 2.0 lanes via the P-links.

All dies (chiplets) interconnect with each other via AMD Infinity Fabric technology. Each CCD connects to the IOD via its own GMI connection. The IOD provides up to twelve Unified Memory Controllers (UMCs) that support DDR5 memory. The IOD also presents 4 'P-links' that the system OEM/ platform designer can configure to support various I/O interfaces, such as PCIe Gen5, and/or CXL 2.0.

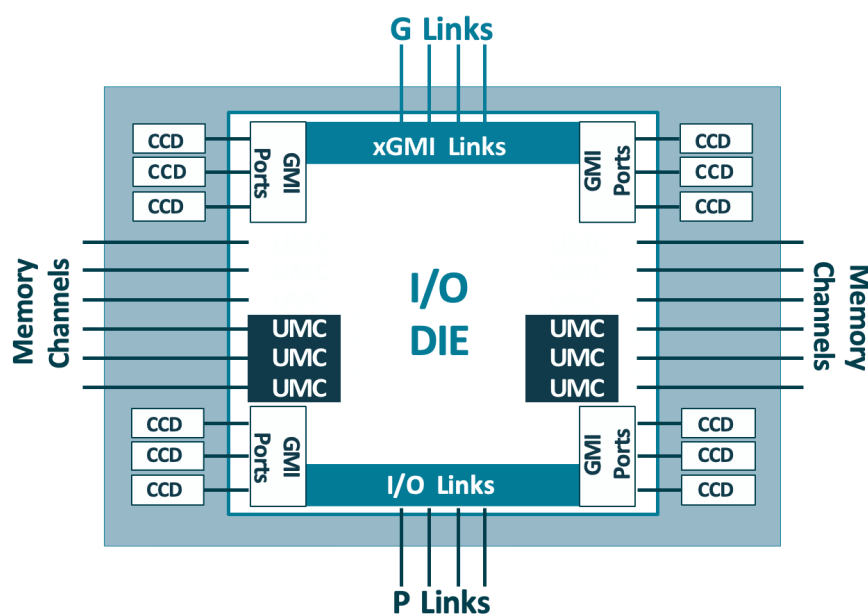


Figure 2-5: AMD EPYC 9005 IOD AMD

AMD also provides “wide” OPNs where each CCD connects to two GMI interfaces, thereby allowing double the Core-to-I/O die bandwidth.

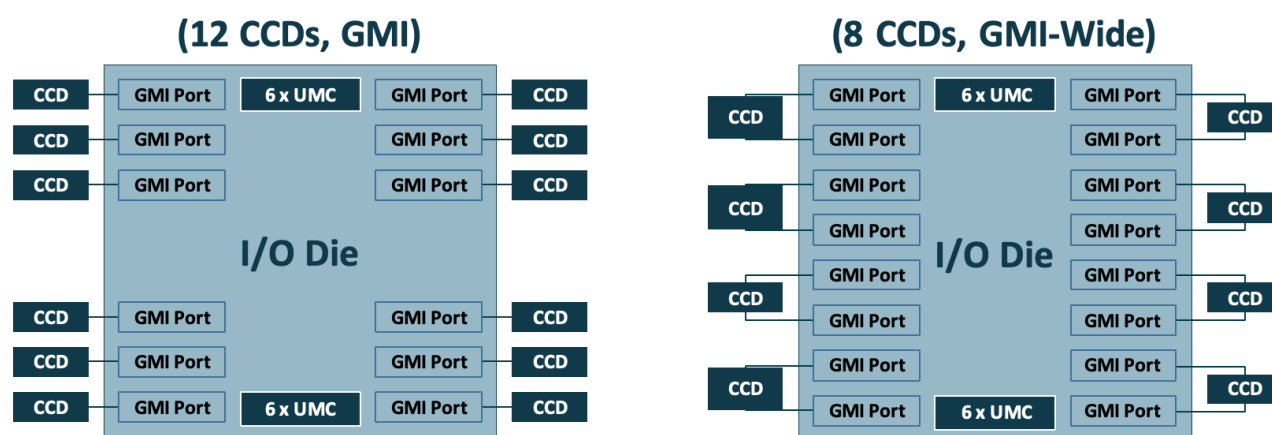


Figure 2-6: Standard GMI links (12 CCD example) vs. Wide GMI links (8 CCD example)

The IOD provides up to twelve Unified Memory Controllers (UMCs) that support DDR5 memory. The IOD also presents 4 ‘P-links’ that the system OEM/platform designer can configure to support various I/O interfaces, such as PCIe Gen5, and/or CXL 2.0.

MEMORY, I/O, AND CONECTIVITY

Each UMC can support up to 2 DIMMs per channel (DPC) for a maximum of 24 DIMMs per socket. OEM server configurations may allow either 1 DIMM per channel or 2 DIMMs per channel. 5th Gen AMD EPYC processors can support up to 9 TB of DDR5 memory. Having additional and faster memory channels compared to previous generations of AMD EPYC processors provides additional memory bandwidth to feed high-core-count processors.

2DPC configurations may provide maximum memory capacity, while 1DPC configurations may provide lower latency. Memory interleaving across channels helps optimize for a variety of workloads and memory configurations.

Each processor may have a set of 4 P-links and 4 G-links. An OEM motherboard design can use a G-link to either connect to a second 5th Gen AMD EPYC processor or to provide additional PCIe Gen5 lanes. A single 5th Gen AMD EPYC processor supports up to eight sets of x16-bit I/O lanes, that is, 128 lanes of high-speed PCIe Gen5 in single-socket platforms and up to 160 lanes in dual-socket platforms. Further, OEMs may either configure 32 of these 128 lanes as SATA lanes or configure 64 lanes as CXL 2.0. In summary, these processors can support:

- Up to 4 G-links of AMD Infinity Fabric connectivity for 2P designs.
- Up to 8 x16-bit (128 lanes) of PCIe Gen 5 connectivity to peripherals in 1P designs (and up to 160 lanes in 2-socket designs).
- Up to 64 lanes (4 P-links) that can be dedicated to Compute Express Link (CXL) 2.0 connectivity to extended memory.
- Up to 32 I/O lanes that can be configured as SATA disk controllers.

CHAPTER 3: NUMA TOPOLOGY

AMD EPYC 9005 Series Processors use a Non-Uniform Memory Access (NUMA) architecture where different latencies may exist depending on the proximity of a processor core to memory and I/O controllers. Using resources within the same NUMA node provides uniform good performance, while using resources in differing nodes increases latencies.

NUMA SETTINGS

A user can adjust the system **NUMA Nodes Per Socket** (NPS) BIOS setting to optimize the NUMA topology for their specific operating environment and workload. For example, setting NPS=4 divides the processor into quadrants, where each quadrant has up to 4 CCDs, 3 UMCs, and corresponding I/O Controller Hub (see Figure 3-1, on the next page). The closest processor-memory I/O distance is between the cores, memory, and I/O peripherals within the same quadrant. The furthest distance is between a core and memory controller or IO hub in cross-diagonal quadrants (or the other processor in a 2P configuration). The locality of cores, memory, and IO hub/devices in a NUMA-based system is an important factor when tuning for performance.

The NPS setting also controls the interleave pattern of the memory channels within the NUMA Node. Each memory channel within a given NUMA node is interleaved. The number of channels interleaved decreases as the NPS setting gets more granular. For example, on a 1P system:

- A setting of NPS=4 partitions the processor into four NUMA nodes per socket with each logical quadrant configured as its own NUMA domain. Memory is interleaved across the memory channels associated with each quadrant. PCIe devices will be local to one of the four processor NUMA domains, depending on the IOD quadrant that has the corresponding PCIe root complex for that device.
- A setting of NPS=2 configures each processor into two NUMA domains that groups half of the cores and half of the memory channels into one NUMA domain, and the remaining cores and memory channels into a second NUMA domain. Memory is interleaved across the six memory channels in each NUMA domain. PCIe devices will be local to one of the two NUMA nodes depending on the half that has the PCIe root complex for that device.
- A setting of NPS=1 indicates a single NUMA node per socket. This setting configures all memory channels on the processor into a single NUMA node. All processor cores, all attached memory, and all PCIe devices connected to the SoC are in that one NUMA node. Memory is interleaved across all memory channels on the processor into a single address space.
- A setting of NPS=0 indicates a single NUMA domain of the entire system (across both sockets in a two-socket configuration). This setting configures all memory channels on the system into a single NUMA node. Memory is interleaved across all memory channels on the system into a single address space. All processor cores across all sockets, all attached memory, and all PCIe devices connected to either processor are in that single NUMA domain.

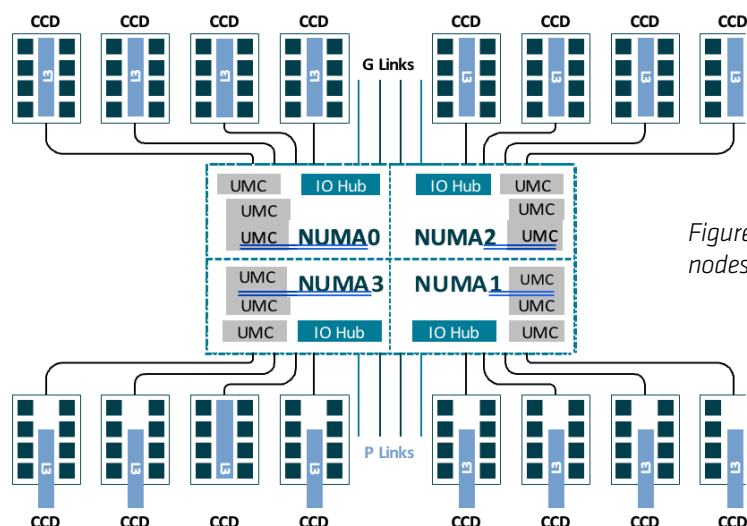


Figure 3-1: 128-core Zen5 processor NUMA nodes with NPS=4

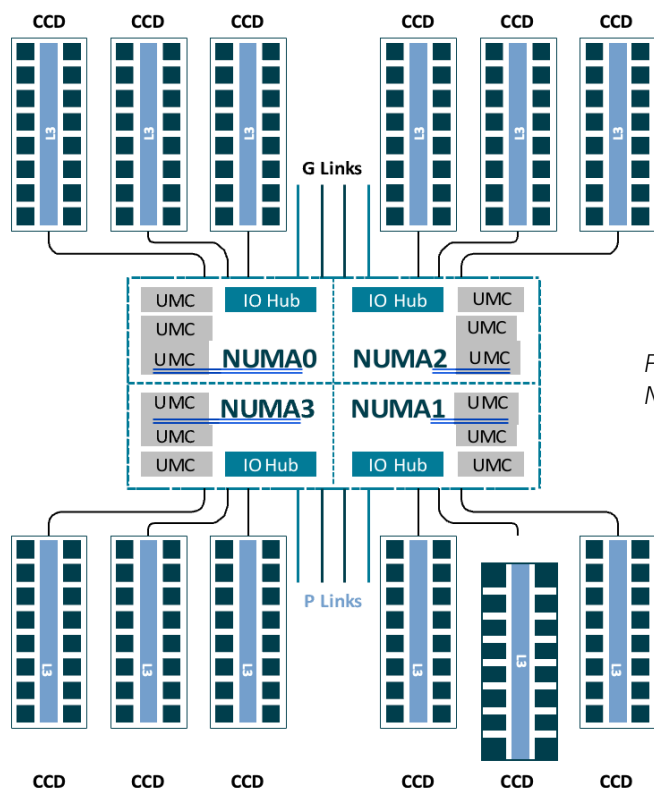


Figure 3-2: 192-core Zen5c processor NUMA nodes with NPS=4

You may also be able to further improve the performance of certain environments by using the LLC (L3 Cache) as NUMA BIOS setting to associate workloads to compute cores that all share a single LLC. Enabling this setting equates each shared L3 or CCX to a separate NUMA node, as a unique L3 cache per CCD. Thus, a single EPYC 9005 Series Processor may support a variety of NUMA configurations ranging from one to sixteen NUMA nodes per socket.

Note: If software needs to understand NUMA topology or core enumeration, it is imperative to use documented Operating System (OS) APIs, well-defined interfaces, and commands. Do not rely on past assumptions about settings such as APICID or CCX ordering.

DUAL-SOCKET CONFIGURATIONS

AMD EPYC 9005 Series Processors support single- or dual-socket system configurations. Processors with a 'P' suffix in their name are optimized for single-socket configurations (see the "Processor Identification" chapter) only. Dual-socket configurations require both processors to be identical. You cannot use two different processor Ordering Part Numbers (OPNs) in a single dual-socket system.

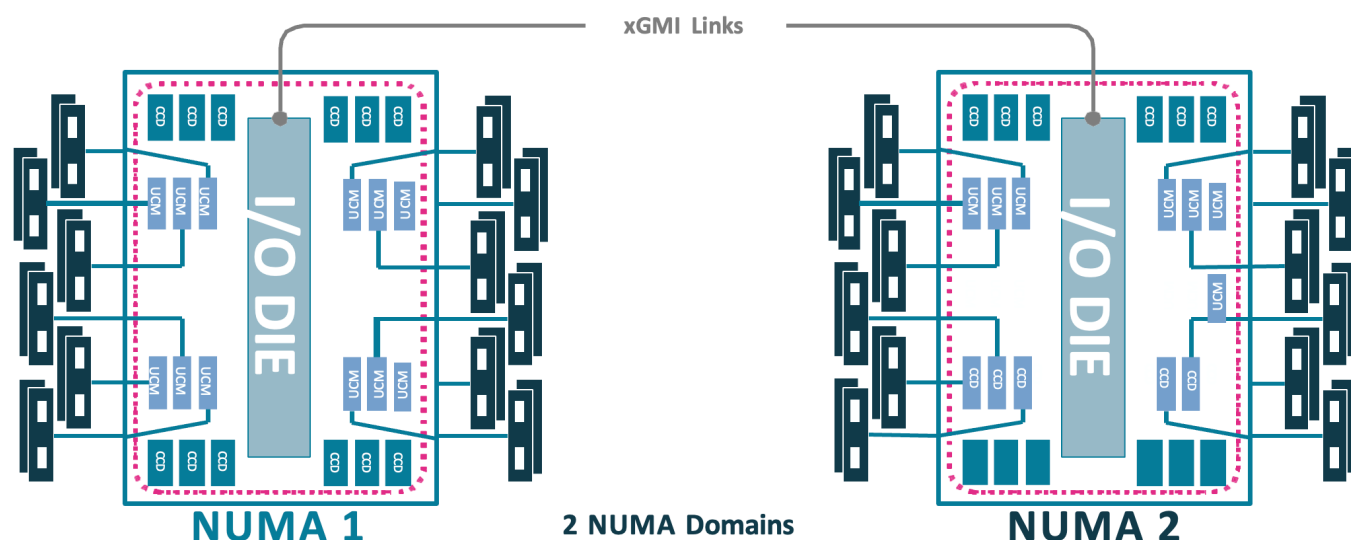


Figure 3-3: Two EPYC 9005 Processors connect through 4 xGMI links (NP51)

In dual-socket systems, two identical EPYC 9005 series SoCs are connected via their corresponding External Global Memory Interconnect [xGMI] links. This creates a high bandwidth, low latency interconnect between the two processors. System manufacturers can elect to use either 3 or 4 of these xGMI/Infinity Fabric links depending upon I/O and bandwidth desires and system design objectives.

The xGMI/Infinity Fabric links utilize the same physical connections as the PCIe lanes on the system. Each link uses up to 16 PCIe lanes. A typical dual socket system will reconfigure 64 PCIe lanes (4 links) from each socket for Infinity Fabric connections. This leaves each socket with 64 remaining PCIe lanes, meaning that the system has a total of 128 PCIe lanes. In some cases, a system designer may want to expose more PCIe lanes for the system by reducing the number of Infinity Fabric G-Links to from 4 to 3. In these cases, the designer may allocate up to 160 lanes for PCIe (80 per socket) by utilizing only 48 lanes per socket for Infinity Fabric links instead of 64.

A dual-socket system has a total of 24 memory channels, or 12 per socket. Different OPNs can be configured to support a variety of NUMA domains.

CHAPTER 4: PROCESSOR IDENTIFICATION

Figure 4-1 shows the processor naming convention for AMD EPYC 9005 Series Processors and how to use this convention to identify particular processors models:

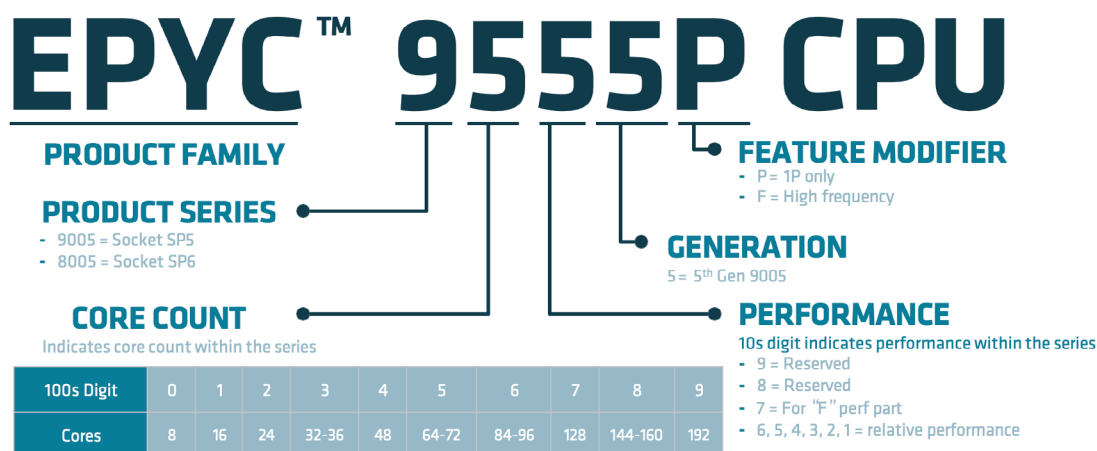


Figure 4-1: AMD EPYC SoC naming convention

CPUID INSTRUCTION

Software uses the CPUID instruction (Fn0000_0001_EAX) to identify the processor and will return the following values:

- **Family:** 1Ah identifies the "Zen 5" architecture.
- **Model:** Varies with product. For example, EPYC Model 10h corresponds to an "A" part "Zen 5" CPU. family
- **Stepping:** May be used to further identify minor design changes.

For example, CPUID values for Family, Model, and Stepping (decimal) of 25, 17, 1 correspond to a "B1" part "Zen 5" CPU.

CHAPTER 5: ADDITIONAL INFORMATION

AMD EPYC 9005 Series Processors introduce several new features that enhance performance. 5th Gen AMD EPYC processors also include ISA updates, additional security features, and improved system reliability and availability compared to prior-gen AMD EPYC processors. Please also see the latest version of the AMD64 Architecture Programmer's Manuals or the *Processor Programming Reference* (available via [https:// devhub.amd.com](https://devhub.amd.com); long required) for the appropriate model and stepping of your AMD EPYC processor.

Not all operating systems or hypervisors support all features. Please refer to your OS or hypervisor documentation for specific releases to identify support for these features.

AVX-512

5th Gen AMD EPYC processors include Advanced Vector Extensions with full 512-bit path and register support (i.e., single instruction, multiple data [SIMD]) operations. 5th Gen AMD EPYC processors also support the new VP2INTERSECT instruction.

CHAPTER 6: RESOURCES

Please see the following resources for additional information about AMD EPYC 9005 Series processors:

- **AMD EPYC™ 9005 Series Server Processors**
- **AMD64 Architecture Programmer's Manual**
- **AMD Documentation Hub**
- *BIOS & Workload Tuning Guide for AMD EPYC™ 9005 Series Processors* (available from the **AMD Documentation Hub**)
- **Memory Population Guidelines for AMD Family 1Ah Models 00h–0Fh and Models 10h–1Fh Socket SP5 Processors** - Login required; please review the latest version if multiple versions are present.
- **Socket SP5/SP6 Platform NUMA Topology for AMD Family 1Ah Models 00h–0Fh and Models 10h–1Fh** - Login required; please review the latest version if multiple versions are present.

DISCLAIMERS

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 particular 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.

COPYRIGHT NOTICE

© 2025 AMD, the AMD Arrow logo, AMD EPYC, 3D V-Cache, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names and links to external sites used in this publication are for identification purposes only and may be trademarks of their respective companies.