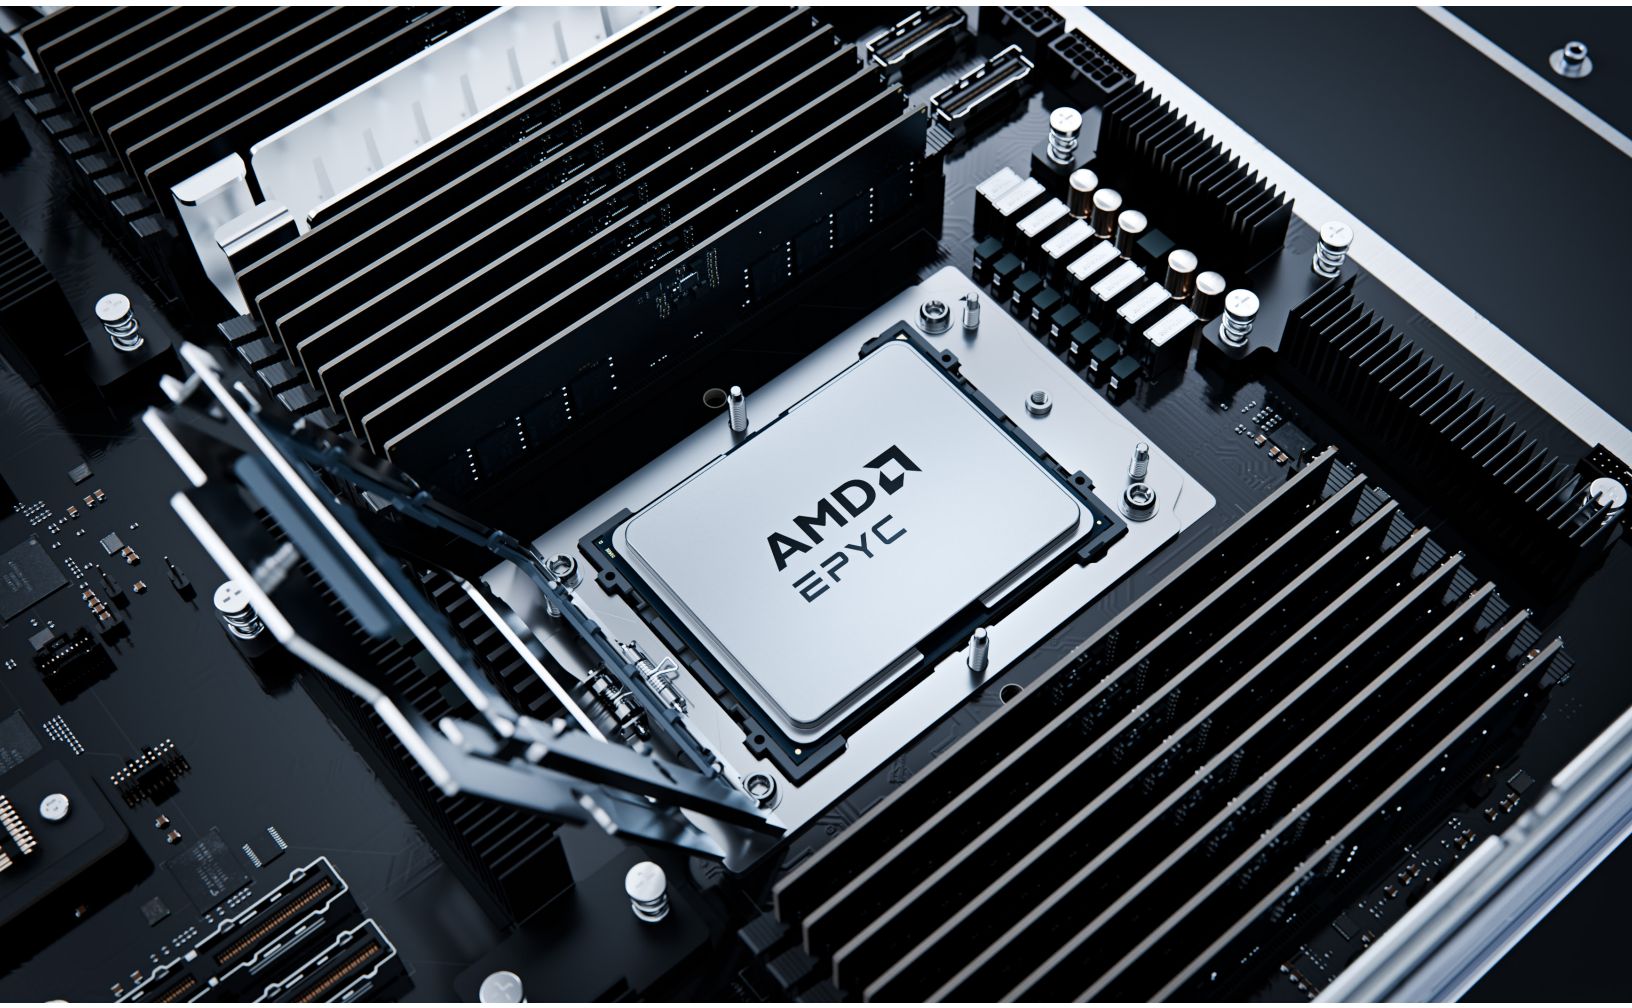


SCALING MEMORY BANDWIDTH ON AMD EPYC™ 9006 SERVER CPUs

DDR5-8000 AND MRDIMM-12800 – SERVER MEMORY TECHNOLOGY, PERFORMANCE, AND WORKLOAD GUIDANCE

TABLE OF CONTENTS

EXECUTIVE SUMMARY **3**

1. INTRODUCTION: THE COMPUTE–MEMORY BALANCE IN MODERN SERVERS **4**

1.1 WHY MEMORY BANDWIDTH PER SOCKET MATTERS 4

1.2 THE MEMORY TECHNOLOGY LANDSCAPE 4

1.3 THE TWO MEMORY OPTIONS ON AMD EPYC 9006 SERVER CPUS 5

2. MRDIMM TECHNICAL DEEP DIVE **5**

2.1 JEDEC CONTEXT AND STANDARDIZATION 6

2.2 INTERNAL ARCHITECTURE (CONCEPTUAL) 6

2.3 VENDOR ECOSYSTEM AND IMPLEMENTATION VARIABILITY 7

2.4 FIRMWARE, BIOS, AND PLATFORM CONSIDERATIONS 7

3. PERFORMANCE CHARACTERISTICS AND EXPECTED BEHAVIORS **7**

3.1 BANDWIDTH: THE CORE BENEFIT 7

3.2 LATENCY AND THE MULTIPLEXING EFFECT 7

3.3 LOAD SCENARIOS 7

3.4 WORKLOAD SENSITIVITY 7

4. RELIABILITY AND SECURITY **8**

4.1 RELIABILITY (RAS AND ECC) 8

4.2 SECURITY 8

5. USE CASES **8**

5.1 CORE USE CASES 8

5.2 VALUE PROPOSITIONS 8

5.3 ADDRESSING POTENTIAL CHALLENGES 9

6. COMPETITIVE AND MARKET CONTEXT **9**

6.1 MARKET CONTEXT 9

6.2 PLATFORM CAPABILITY & TIMING 9

7. CONCLUSION **9**

APPENDICES **10**

APPENDIX A – GLOSSARY 10

APPENDIX B – REFERENCES 10

EXECUTIVE SUMMARY

Memory bandwidth has become a frequent constraint for modern server workloads. As core counts rise, performance increasingly depends on how quickly memory can feed the cores: HPC solvers, CPU-based and agentic AI inference, data preprocessing that feeds accelerators, and large in-memory analytics are all often bound by memory bandwidth. AMD has consistently expanded memory bandwidth alongside compute capacity, and the AMD EPYC 9006 Server CPUs continue that commitment with a DDR5 memory subsystem built to keep high-core-count processors supplied with data.

The portfolio offers two high-performance memory options: native DDR5-8000 RDIMMs for broad deployment, and JEDEC-standard second generation MRDIMMs (Gen2) at up to 12,800 MT/s for the most bandwidth-intensive workloads – delivering up to 1.6 TB/s of memory bandwidth per socket on the 16-channel platform.[2] In the agentic AI era, this bandwidth lets the AMD EPYC 9006 Server CPUs scale more agents, maximize AI host-node throughput, and keep the GPUs they feed supplied with data.

DDR5 RDIMM technology is a well understood and widely deployed memory technology defined by the JEDEC® industry organization that develops the critical standards for memory safety and interoperability. A newer differentiator in the market is the advent of JEDEC-standard MRDIMM (Multiplexed Rank DIMM) technology. But even these are not wholly "new" technologies but are a new implementation of proven memory module technologies. These new MRDIMMs operate two ranks of DRAM in parallel and multiplex them through an on-module data buffer. As such, these modules present the memory controller roughly double the per-channel data rate of a same-generation RDIMM – while keeping the same protocol-compatible DDR5 interface the platform already understands.

This paper explains how MRDIMM works, why it matters on the AMD EPYC 9006 Server CPUs, and how to evaluate it for memory-bound workloads along with manageability, reliability, and security considerations.

1. INTRODUCTION: THE COMPUTE–MEMORY BALANCE IN MODERN SERVERS

1.1 WHY MEMORY BANDWIDTH PER SOCKET MATTERS

Across recent server generations, per-socket core counts have grown faster than per-socket memory bandwidth, so the bandwidth available per core has generally trended flat-to-down, creating potential challenges to scalability. As AMD EPYC 9006 Server CPUs feature up to 256 cores / 512 threads per socket, the memory system must quickly deliver more data or those cores could be "starved" of data.[1] The metrics to consider are not capacity alone but bandwidth-per-core: a bandwidth-starved core cannot do useful work waiting on data. Increasing usable bandwidth per socket raises throughput per server, improves CPU utilization by keeping cores fed rather than waiting on data and can reduce node count required to accomplish a fixed amount of work.

The long-run trend is not that absolute memory bandwidth has stalled – it has grown enormously thanks to CPU and memory technology advancements – but the fact is that core growth has outpaced it. Memory DIMM / DDR-generation improvements (3200 MT/s DDR4 vs 6400 MT/s DDR5) and CPU memory channel increases (EPYC 9004 family with 12-channel DDR5, EPYC 9006 family with up to 16-channels) help balance generational processor core count gains. The advent of JEDEC-standard MRDIMM-12800 (Gen2) technology will provide a major boost to help balance CPU and memory bandwidth needs:

EPYC SERVER CPU GENERATION	MAX CORES (UP TO)	MEMORY TECHNOLOGY	~PEAK BW/SOCKET	~BANDWIDTH/CORE
EPYC 7001 (2017)	32	8ch RDIMM DDR4-2666	~170 GB/s	~5.3 GB/s
EPYC 7002/7003 (2019–21)	64	8ch RDIMM DDR4-3200	~205 GB/s	~3.2 GB/s
EPYC 9004 (2022)	96	12ch RDIMM DDR5-4800	~460 GB/s	~4.8 GB/s
EPYC 9005 (2024)	192	12ch RDIMM DDR5-6400	~614 GB/s	~3.2 GB/s
EPYC 9006 (2026)	256[1]	16ch RDIMM DDR5-8000	~1.0 TB/s[5]	~4.0GB/s[4]
EPYC 9006 + MRDIMM-12800 (2026)	256	16ch MRDIMM-12800	~1.6 TB/s[2]	~6.4 GB/s[4]

Capacity and bandwidth are distinct needs – large in-memory databases, dense virtualization, and big working data sets can need terabytes of addressable DRAM. But capacity and bandwidth are distinct considerations with distinct solutions, and conflating them can lead to a suboptimal DIMM or system configuration choice and unnecessary expense.

1.2 THE MEMORY TECHNOLOGY LANDSCAPE

RDIMM (Registered DIMM) – RDIMMs are the general-purpose server memory technology in ubiquitous use today; on AMD EPYC 9006 Server CPU systems they support DDR5-6400 and DDR5-8000 and reach the highest per-module capacities on the market. An RDIMM buffers command/address signals through a registering clock driver (RCD) and provides on-module Error-Correcting Code (ECC) technology. The data path runs at the native DRAM data rate. As memory population rises (2 DIMMs per channel, higher-density modules), electrical loading can constrain the top stable data rate – a signal-integrity limit, addressed at the platform and module level.

MRDIMM (Multiplexed Rank DIMM) – The AMD EPYC 9006 Server CPUs are the first AMD server platforms to support the 2nd Generation MRDIMM technology. MRDIMM Gen2 is a JEDEC-defined DDR5 module class that increases per-channel memory bandwidth. It reads two ranks of DRAM in parallel and time-multiplexes their data onto the memory channel, using standardized on-module logic – a set of Multiplexed-rank Data Buffers (MDBs) and a Multiplexed-rank Registering Clock Driver (MRCD). This presents the host with roughly double the effective data rate, while the DRAM runs at about half that rate. Crucially, the interface to the memory controller remains protocol-compatible DDR5; the multiplexing is handled on-module.

LPDDR – Low-power DDR (LPDDR) is an alternative memory now being used on selected systems. While this technology has been available in smartphones and mobile devices for years, these JEDEC-standard memories are now also available in field replaceable units called Small Outline Compression Attached Memory Module (SOCAMM2) modules. The primary goal of this memory is leadership energy efficiency (performance-per-watt). On a per-module basis, through a wider interface and the higher per-pin data rates of LPDDR5X, it can deliver roughly twice the bandwidth of a standard RDIMM at significant lower power, though with lower capacity – an advantage for power- and space-constrained designs. LPDDR comes in the SOCAMM2 form factor rather than the familiar DIMM format. SOCAMM2 modules lie flat on system boards and have a different socket type than traditional DIMMs, so they require purpose-built system designs and are not compatible with RDIMM systems. To date, this type of memory is not widely used on mainstream server products; with selected SKUs of AMD EPYC 9006 Server CPUs (AMD EPYC 9006 LP) based systems providing LPDDR5X support in 2027.

HBM – There are a number of JEDEC standards for "High Bandwidth Memory" HBM technologies. These cover a range of performance and cost points, but in general, HBM delivers very high bandwidth through stacked DRAM co-packaged with the base logic die and it is comparatively expensive due to the premium performance and the significant manufacturing complexities in building stacked memories. It is included here only to position it relative to DIMM-based memory; it is outside the scope of this paper.

1.3 THE TWO MEMORY OPTIONS ON AMD EPYC 9006 SERVER CPUs

DDR5 brings higher DRAM device densities (24 Gb and 32 Gb devices), improved channel/bank structure, and on-module power management compared to prior generation DDR4. On AMD EPYC 9006, the two high-performance memory options are:

- **DDR5-8000 RDIMM** – Registered DIMM running its DRAM natively at 8,000 MT/s. Broad availability, standard capacities, will be the default for most deployments with a good balance of performance and commodity-like cost structures.
- **MRDIMM-12800 (Gen2)** – Multiplexed Rank DIMM presenting 12,800 MT/s to the host by multiplexing two ranks of ~6,400 MT/s DRAM. Its higher component count typically makes it a premium offering, but it is the higher-bandwidth option for memory-bound workloads.

Platform note. The bandwidth figures in this paper refer to the 16-channel SP7 platform, the maximum-bandwidth option. AMD EPYC 9006 is also offered on the 8-channel SP8 platform, which supports MRDIMM-12800 (one DIMM per channel) but delivers proportionally lower per-socket bandwidth; size SP8 deployments against its channel count rather than the SP7 figures used here.

2. MRDIMM TECHNICAL DEEP DIVE

- **Bandwidth is the primary benefit.** MRDIMM's advantage is per-channel and per-socket bandwidth, and at a comparable configuration it is expected to be available at equivalent capacities, though 3DS RDIMM reaches larger per-module sizes. But with premium performance and a higher bill of materials for any given capacity point, it is expected to be more costly on a $/GB basis than RDIMMs. To determine the best value solution for a deployment, match the DIMM type to whether the workload is bandwidth-
or capacity-constrained.
- **Underlying DRAM runs slower than the interface.** At MRDIMM-12800 the host sees 12,800 MT/s while the DRAM devices each run at ~6,400 MT/s. Running the individual memory devices slower is what keeps per-device timing margins and power in check. Under load, latency can also improve rather than degrade because the added bandwidth keeps the channel further from saturation (see §3.2). At the module level, though, the faster channel and on-module buffers mean an MRDIMM draws somewhat more power than a comparable RDIMM – a factor in thermal and airflow planning.
- **Standards-based, broad-ecosystem.** The MDB and MRCD logic are published JEDEC standards, so the 2nd generation of MRDIMM is a multi-vendor ecosystem (Micron, Samsung, SK hynix). Vendor modules still differ in firmware/telemetry and should be validated per platform. Note that Intel and a select set of memory partners were offering up to 8,800MT/s capable "MRDIMMs" for specific 6th Generation Intel® Xeon® SKUs, but this is a pre-standard, first-generation technology developed before the JEDEC MRDIMM standard was finalized.
- **Protocol-compatible interface minimizes design change.** The system talks DDR5 physically and electrically; the multiplexing is on-module. Platform enablement (training, thermals, firmware) is required – see §4 and §7, but both RDIMM and MRDIMM use the same socket/connectors.

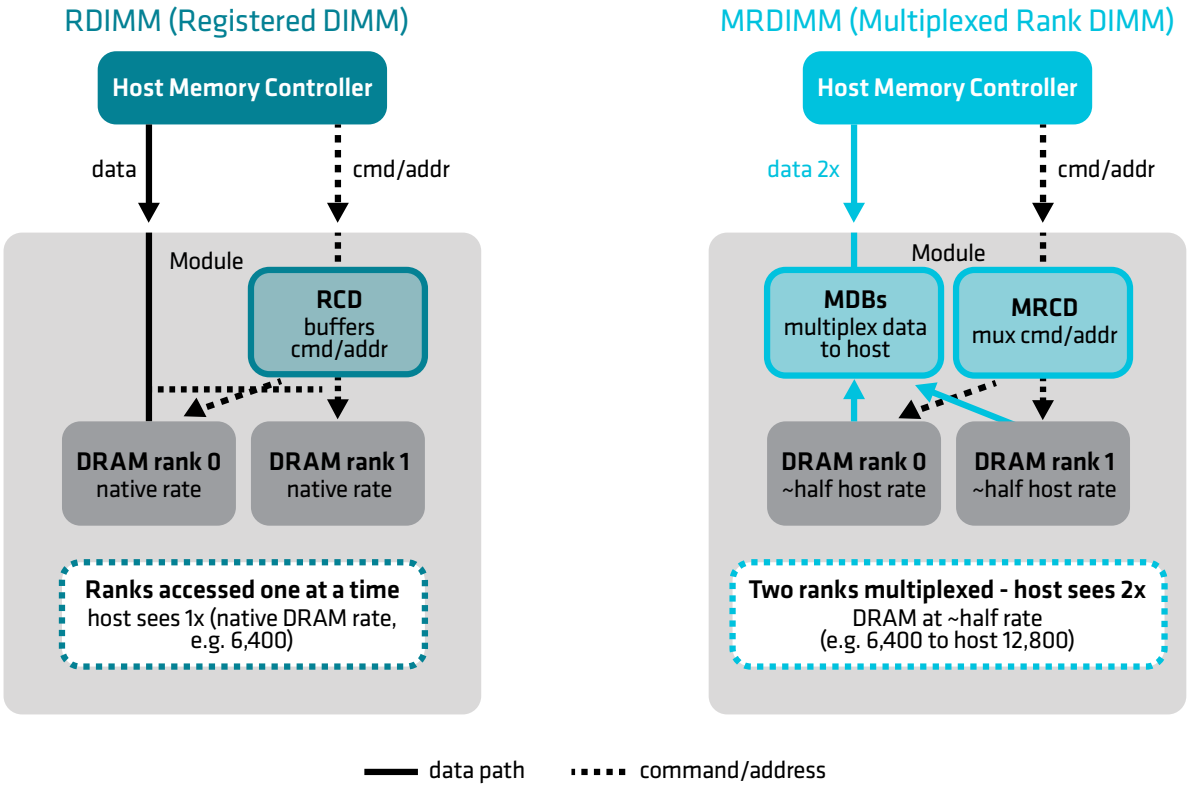
2.1 JEDEC CONTEXT AND STANDARDIZATION

MRDIMM is standardized by JEDEC as part of the DDR5 module ecosystem. The relevant logic standards are the Multiplexed-rank Data Buffer (MDB) and the Multiplexed-rank Registering Clock Driver (MRCD); JEDEC published the Gen2 MDB logic standard (JESD82-552 / DDR5MDB02) in 2026, with the corresponding MRCD standard following. The JEDEC roadmap is generational, with the Gen2 logic supporting 12,800 MT/s and a Gen3 target of even higher data rate.

2.2 INTERNAL ARCHITECTURE (CONCEPTUAL)

- **Rank multiplexing.** Two ranks operate in parallel; the MDB multiplexes their data and drives the channel at double the per-rank rate, absorbing the electrical load that would otherwise cap signaling speed.
- **Command/address path.** The MRCD handles timing for the multiplexed ranks, analogous to an RCD but multiplex-aware.
- **Management path.** Modules expose SPD/EEPROM and may add vendor telemetry/firmware for health and updates (see §4).

Figure 1 - RDIMM vs. MRDIMM Architecture



MRDIMM keeps a protocol-compatible DDR5 interface; the multiplexing (MDBs + MRCD) is handled on-module. JEDEC-standard logic.

2.3 VENDOR ECOSYSTEM AND IMPLEMENTATION VARIABILITY

All leading memory vendors have public MRDIMM roadmaps and offer DDR5-8000 RDIMM and DDR5-12800 MRDIMM products. A tall form factor (TFF) MRDIMM variant extends module height to fit more DRAM packages, reaching higher capacity (~256 GB with single-die devices avoiding costlier 3DS/TSV stacking) and improving thermals, at the cost of chassis clearance (typically 2U+). Firmware, telemetry, and exact latency behavior differ by vendor; module behavior should be validated on the target AMD EPYC 9006 Server CPU platform.

2.4 FIRMWARE, BIOS, AND PLATFORM CONSIDERATIONS

MRDIMM requires platform enablement: BIOS/UEFI support, correct memory training/initialization for multiplexed ranks, and validated thermal/airflow provisioning (MRDIMM draws more power per module than RDIMM.) Use the OEM's validated BIOS revision and a supported MRDIMM configuration.

3. PERFORMANCE CHARACTERISTICS AND EXPECTED BEHAVIORS

3.1 BANDWIDTH: THE CORE BENEFIT

Per channel, the math is direct. DDR5-6400 provides ~51.2 GB/s of raw memory bandwidth; MRDIMM-12800 (Gen2) ~102.4 GB/s – breaking the 100 GB/s-per-channel barrier and roughly doubling DDR5-6400. Across 16 channels, that is how AMD EPYC 9006 Server CPUs reach up to ~1.6 TB/s per socket. DDR5-8000 RDIMM provides ~64 GB/s per channel – the strong default when bandwidth demand is moderate or capacity or cost are the higher priorities over absolute performance.[2]

MODULE (PER CHANNEL)	HOST RATE	APPROX. PEAK BANDWIDTH PER CHANNEL
DDR5-6400 RDIMM (baseline ref)	6,400 MT/s	~51.2 GB/s
DDR5-8000 RDIMM	8,000 MT/s	~64 GB/s
MRDIMM-12800 (Gen2)	12,800 MT/s	~102.4 GB/s (>100; ~2× 6400)

3.2 LATENCY AND THE MULTIPLEXING EFFECT

Under bandwidth pressure, MRDIMM can reduce loaded memory-access latency: as a channel approaches saturation, requests spend less time queued because the channel carries roughly twice the data. Independent testing on prior-generation MRDIMM shows this latency benefit grows with bandwidth utilization. This corrects a common misconception: the on-module buffering does not necessarily increase overall latency relative to RDIMM. Unloaded, single-DIMM latency can be broadly comparable to RDIMM.

3.3 LOAD SCENARIOS

- **Light load:** MRDIMM adds only minimal latency on the unloaded path from the multiplexing, so unloaded latency is close to RDIMM.
- **Heavy load/bandwidth pressure:** MRDIMM sustains higher effective bandwidth and offers better loaded latency than RDIMM at a given bandwidth point.
- **Capacity-led configurations: for the highest per-socket capacity:** RDIMM reaches the largest modules; MRDIMM leads on GB/s. Choose per the workload's binding constraint.

3.4 WORKLOAD SENSITIVITY

Most improved by MRDIMM: Memory-bandwidth-bound work –sparse-matrix HPC solvers (as measured by the HPCG benchmark) , CPU-based AI inference, data preprocessing feeding GPUs/accelerators, large in-memory analytics. **Least sensitive:** compute-bound work, or capacity-bound work where the constraint is GB, not GB/s (use high-capacity RDIMM there).

4. RELIABILITY AND SECURITY

MRDIMM preserves the same reliability and security characteristics as RDIMM. It is a standards-based DDR5 module class that keeps the established RAS feature set and platform-level protections – the on-module multiplexing adds bandwidth, not a reliability or security trade-off.

4.1 RELIABILITY (RAS AND ECC)

MRDIMM supports the full DDR5 RAS feature set – the same as RDIMM. This includes side-band ECC (using ECC DRAM devices) and on-die ECC, error check-and-scrub (ECS), and post-package repair (PPR). The on-module data buffers (MDB) carry ECC through the multiplexing path, so end-to-end error protection is maintained across the two multiplexed ranks. In addition, AMD EPYC 9006 Server CPUs support link CRC protection on the Data and Command/Address pins for MRDIMM, in line with the higher data rate. Correctable and uncorrectable errors are logged and surfaced through standard interfaces (BMC, OS, Redfish), and module health telemetry integrates with existing server monitoring. The one practical difference from RDIMM is the additional on-module logic (MDB/MRCD), which carries its own firmware and telemetry.

4.2 SECURITY

Security is foundational to the memory subsystem design in AMD EPYC 9006 Server CPUs. Both memory technologies are protected by AMD Infinity Guard, a layered security foundation spanning silicon-rooted trust, secure boot, attestation, and confidential computing (AMD SEV).[6] Transparent Secure Memory Encryption (TSME), which transparently encrypts DRAM regardless of DIMM type; MRDIMM introduces no change to this protection. Because TSME operates in the processor's memory controller, it applies uniformly to RDIMM and MRDIMM alike. At the module level, MRDIMM supports signed firmware and secure update mechanisms for the on-module logic.

5. USE CASES

5.1 CORE USE CASES

WORKLOAD	WHY BANDWIDTH IS THE LIMITER	RECOMMENDED
General-purpose enterprise	Balanced; capacity and cost concerns usually dominate	DDR5-8000
Virtualization / consolidation	Often capacity-led; BW matters for heavy VMs	DDR5-8000; MRDIMM for BW-heavy hosts
In-memory DB & analytics	Scan/aggregate are bandwidth-bound; large working sets	MRDIMM or DDR5-8000 by capacity
Data preprocessing / feeding GPUs	Host memory BW gates accelerator utilization	MRDIMM-12800 preferred
Agentic AI & CPU-based LLM inference	Token generation re-reads the KV cache from memory each step – bandwidth-bound	MRDIMM-12800 preferred
HPC / scientific (HPCG-class)	Sparse solvers stall on memory, not FLOPs	MRDIMM-12800 preferred

WHY LLM INFERENCE IS BANDWIDTH-BOUND:

During generation, a large language model produces one token at a time, and for each new token it must re-read a large, growing dataset from memory – the model weights plus the KV cache (the stored key/value vectors of all prior tokens, which grows with context length). Token-generation throughput is therefore gated by how fast memory can be read, not by compute. Higher per-socket bandwidth reads this data faster, yielding more tokens per second – which is why MRDIMM-12800 benefits CPU-based inference and agentic AI workloads running many concurrent agents, where the AMD EPYC 9006 Server CPUs keep the memory-bound generation step fed.

5.2 AMD EPYC 9006 SERVER CPUs MEMORY SUPPORT BENEFITS

- **More work per server on bandwidth-bound workloads** by keeping high-core-count AMD EPYC 9006 Server CPUs fed – better utilization of hardware already bought.
- **Standards-based** JEDEC MRDIMM from multiple suppliers – supply assurance and interoperability.
- **A clean two-option story:** DDR5-8000 for broad deployment and capacity; MRDIMM-12800 when the workload is bandwidth-bound.

5.3 ADDRESSING POTENTIAL CHALLENGES

A few practical considerations shape the MRDIMM decision. On cost, MRDIMM is a premium, higher-component-count module and carries a price premium over RDIMM; it is best suited to workloads where the bandwidth gain justifies the added cost, evaluated per deployment on a performance-per-dollar and throughput-per-server basis. On latency, under load MRDIMM often improves memory-access latency rather than degrading it – the added bandwidth keeps the channel further from saturation, so requests spend less time queued – while unloaded latency is comparable to RDIMM. On capacity, high-capacity RDIMM remains the better fit for capacity-first designs, with MRDIMM reserved for bandwidth-first deployments. And on operations, because MRDIMM is JEDEC-standard and uses OEM-validated parts in the same DDR5 slots as RDIMM, it integrates into existing server designs without a new memory paradigm.

6. COMPETITIVE AND MARKET CONTEXT

6.1 MARKET CONTEXT

MRDIMM emerged as the industry-standard answer to the memory-bandwidth wall on DDR5, driven by AI, HPC, and dense cloud workloads. JEDEC standardized the MDB/MRCD logic; all three major suppliers build to it.

6.2 PLATFORM CAPABILITY & TIMING

PLATFORM	CHANNELS	MRDIMM	AVAILABILITY
AMD EPYC 9006 (SP7)	16	MRDIMM-12800 (Gen2) JEDEC Standard	Shipping
AMD EPYC 9006 (SP8)	8	MRDIMM-12800 (Gen2) JEDEC Standard (1DPC)	1H 2027
AMD EPYC 9005 (SP5)	12	None	Shipping
Intel Xeon 6 “Granite Rapids”	Up to 12	MRDIMM-8800, pre-standard	Shipping
Intel Xeon 6+ “Clearwater Forest”	12	None	Shipping (2026)
Intel Xeon 7 “Diamond Rapids”	16	MRDIMM-12800 (Gen2) JEDEC Standard	2027

AMD EPYC 9006 Server CPUs can support up to 16 channels and JEDEC-standard MRDIMM Gen2 memory (12,800 MT/s). Intel’s currently available Xeon 6 platform is up to 12-channel, but it uses a pre-standard first-generation MRDIMM at 8,800 MT/s. Intel’s announced 2027 server product is a 16-channel, 12,800 MT/s platform.[3]

7. CONCLUSION

AMD EPYC 9006 Server CPUs give customers standards-based memory options to match bandwidth to workload demands:

- Choose the right memory for the workload. DDR5-8000 RDIMM provides a strong general-purpose option, while MRDIMM-12800 supports workloads that demand higher memory bandwidth.
- Feed more cores with more bandwidth. MRDIMM-12800 delivers twice the per-channel bandwidth of DDR5-6400 and up to 1.6 TB/s per socket, helping keep high-core-count processors productive.[2]
- Support demanding AI and enterprise workloads. Higher bandwidth helps address bottlenecks in memory-bandwidth-bound HPC, agentic AI, CPU-based inference, and large in-memory analytics, with the potential for lower latency under load.
- Build on standards and supplier choice. A common JEDEC standard and multiple memory suppliers give customers flexibility in sourcing.
- Use familiar DDR5 connections. MRDIMM uses the same DDR5 slots and connectors on platforms designed to support it.

By scaling memory bandwidth alongside compute, AMD EPYC 9006 Server CPUs keep high-core-count processors supplied with data across the workloads that define the modern data center.

APPENDICES

APPENDIX A – GLOSSARY

- **MRDIMM** – Multiplexed Rank DIMM. Multiplexes two ranks to ~2× channel bandwidth.
- **RDIMM** – Registered DIMM. Buffered command/address; data at native DRAM rate.
- **MDB / MRCD** – Multiplexed-rank Data Buffer / Registering Clock Driver (JEDEC logic).
- **TFF** – Tall Form Factor MRDIMM (more DRAM packages; higher capacity/better thermals; needs chassis clearance).
- **ECS / PPR** – Error Check and Scrub / Post-Package Repair. DDR5 RAS features: ECS auto-corrects and rewrites bit errors in the background; PPR remaps a failing DRAM row to a spare.
- **MT/s, GB/s, SI, SPD, ECC, memory bandwidth per core –** standard memory terms.

APPENDIX B – REFERENCES

TYPE	REFERENCE	LOCATOR
JEDEC	DDR5 MRDIMM ecosystem / MDB & MRCD logic standards; expanded roadmap	jedec.org
AMD	AMD EPYC 9006 disclosures (AMD blog + product docs)	amd.com
Suppliers	Micron / Samsung / SK hynix MRDIMM & DDR5-8000 briefs (Computex 2026)	vendor product pages
Intel	Xeon 6 (pre-standard MRDIMM) and Xeon 7 "Diamond Rapids" (2027) disclosures	intel.com
Independent	Independent MRDIMM evaluations on Intel Xeon 6 (Phoronix; arXiv energy/latency study) – prior-generation reference	phoronix.com / arxiv.org
Analyst/Press	TrendForce; ServeTheHome; Tom's Hardware (corroboration only)	various

ENDNOTES AND LEGAL

1. **EPYC-025E:** As of November 2026, 6th Gen AMD EPYC 9996 has 256 cores and 512 threads per socket with SMT enabled, which is higher than any other publicly disclosed server CPU. NVIDIA Vera has 88 cores and 176 threads per socket. Intel Xeon 6980P has 128 cores and 256 threads per socket.
2. **EPYC-032E:** AMD EPYC™ 9006 processors support up to 16 channels of 12,800 MT/s MRDIMM memory, providing up to 1,638.4 GB/s of theoretical memory bandwidth per socket. Intel® Xeon® 6 processors and ARM AGI processors support up to 12 channels of 8,800 MT/s MRDIMM memory, providing up to 844.8 GB/s of theoretical memory bandwidth per socket. AMD EPYC 9006 Series processors are the first CPUs to support JEDEC-standard MRDIMM memory at 12,800 MT/s, as of 11/05/2026.
3. **EPYC-033E:** AMD EPYC 9006 CPUs support 16 memory channels. 6th Gen Intel Xeon CPUs support 12 memory channels. 3rd, 4th, 5th, and some 6th gen Intel Xeon CPUs support 8 memory channels, Intel Xeon Memory specifications per https://ark.intel.com. ARM AGI CPU supports 12 memory channels per https://www.arm.com/products/cloud-datacenter/arm-agi-cpu
4. **9xx6-101:** The AMD EPYC™ 9996 processor with 256 physical cores per socket and 16 memory channels provides up to 1,638.4 GB/s of theoretical memory bandwidth per socket with MRDIMM-12800 memory at 12,800 MT/s (1 DIMM per channel), or 6.4 GB/s per core. With DDR5-8000 RDIMM memory at 8,000 MT/s (1 DIMM per channel), it provides up to 1,024 GB/s per socket, or 4.0 GB/s per core. Theoretical bandwidth as of 11/05/2026; actual bandwidth varies by workload and configuration.
5. **9xx6-102:** AMD EPYC™ 9006 Series processors support up to 16 channels of DDR5-8000 RDIMM memory at 8,000 MT/s (1 DIMM per channel), providing up to 1,024 GB/s (approximately 1.0 TB/s) of theoretical memory bandwidth per socket. Theoretical bandwidth as of 11/05/2026; actual bandwidth varies by workload and configuration.
6. **GD-183A:** AMD Infinity Guard features vary by EPYC™ Processor generations and/or series. 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 **http://www.amd.com/en/products/processors/server/epyc/infinity-guard.html**.