

ESSENTIAL GUIDE TO MICRON DDR5

INTRODUCTION

Development of next-generation memory technology is being driven by the growth of cloud and enterprise data center applications, artificial intelligence, autonomous cars, augmented reality, embedded vision, and High Performance Computing (HPC).

Advanced server workloads resulting from these compute-intensive applications have fueled processor core counts, which are increasing at a rate that surpasses that of total system memory bandwidth available. So in many ways, the key to enabling greater data center workloads is higher-performance, denser, higher-quality memory. DDR5 offers nearly double the effective bandwidth when compared to its predecessor DDR4, helping relieve this bandwidth per core obstacle.

This white paper discusses the benefits of DDR5, challenges of transition, and why migrating to this latest standard not only makes sense but also could well be a competitive difference maker.




EVOLUTION OF DDR STANDARDS

Today, Dual Data Rate (DDR) Synchronous Dynamic Random-Access Memory (SDRAM) technology is the de facto memory used in almost all applications, from data-center to mobile use cases. This is due to DDR's high density and straightforward architecture (using a capacitor as a storage element), low latency, and low-power consumption.

DDR SDRAM achieves double-data bandwidth without increasing the clock frequency by transferring data on both rising and falling edges of the clock signal. In DDR, prefetch buffer size is $2n$ (two data words per memory access), which is double SDR SDRAM's prefetch buffer size. Prefetch, by way of review, allows a single address request to result in multiple data words. That is, when a memory access occurs to a row, a set of adjacent data words will be grabbed and read out in rapid-fire sequence on the I/O pins, without the need for individual column address requests.

DDR2 transferred data at twice the clock speed. It does so by transferring data on the rising and falling edges of the clock signal. What is more, the internal clock runs at half the speed of the data bus, leading to higher bus speed and lower power. Taken together, DDR2-SDRAM achieves four data transfers per internal clock cycle.

Micron's DDR5 Delivering Next Gen Performance



© 2020 Micron Technology, Inc. Micron, the Micron orbit logo, the M orbit logo, Intelligence Accelerated™, and other Micron trademarks are the property of Micron Technology, Inc. All other trademarks are the property of their respective owners.

NOTE: This infographic contains forward-looking statements regarding the production of DDR5. Actual events or results may differ materially from those contained in the forward-looking statements. Please refer to the documents Micron files on a consolidated basis from time to time with the Securities and Exchange Commission, specifically Micron's most recent Form 10-K and Form 10-Q. These documents contain and identify important factors that could cause the actual results for Micron on a consolidated basis to differ materially from those contained in our forward-looking statements (see Certain Factors). Although we believe that the expectations reflected in the forward-looking statements are reasonable, we cannot guarantee future results, levels of activity, performance or achievements.



Figure 1: Micron has been a major participant in each new generation of DDR memory. (Source: Micron)




DDR3 transfers data at twice the rate of DDR2 SDRAM, enabling higher bandwidth and peak data rates. Since the data retention characteristics of DRAM cells depend on temperatures, new features such as Self Refresh Temperature Range and Automatic Self-Refresh allow memory to control the refresh rates according to the temperature variation. When the temperature is high, the self-refresh operation will execute in a short period to prevent data loss. The width of the DDR3 prefetch buffer (a memory cache that stores data before it is actually needed) is 8 bit.

DDR4 kept the prefetch buffer size the same as DDR3 but was able to achieve even higher speed and efficiency by sending more read/write commands per second. DDR4 standard divides the DRAM banks into two or four selectable bank groups, where transfers to different bank groups can be done faster. Operating voltage of DDR4 is reduced to 1.2 V compared to DDR3. New features enhanced DDR4 memory's signal integrity and improved the stability of data transmission/access.

With **DDR5**, the DRAM and buffer chip registering clock driver (RCD) voltage drops from 1.2 V down to 1.1 V, some 8% lower than what DDR4 achieved. On its own, this may not seem like a great deal, but keep in mind companies employ tens of thousands of machines. In tightly packed servers, memory modules can consume hundreds of watts, so it adds up. DDR5 also enables the increased reliability, availability, and serviceability (RAS) that modern data centers require.

Standard DDR DRAMs are available in various types and form factors such as discrete DRAMs or DIMMs. DIMMs are printed-circuit board (PCB) modules with several DRAM chips supporting either a 64-bit or a 72-bit data width. The latter are called error-correcting code (ECC) DIMMs since they have 8 bits of storage for on-die ECC, improving reliability by protecting the memory array against single-bit errors.



With DDR5 DIMMs, power management moves from the motherboard to the DIMM itself; DDR5 DIMMs have their own voltage regulators and 12-V power management integrated circuits (PMICs). The PMIC distributes the 1.1-V VDD supply, helping with signal integrity and noise. These buses already feature on-die termination—a technology where the termination resistor for impedance matching in transmission lines is located inside a semiconductor chip instead of on a printed-circuit board—to make signals cleaner and to improve stability at high data rates.




Figure 2: DDR5 will enable the next generation of server workloads by delivering more than an 85% increase in memory performance. (Source: Micron)

With DDR5, each DIMM will have two channels. Each of these channels is 40-bits wide: 32 data bits with eight ECC bits. While the data width is the same (64-bits total), having two smaller independent channels improves memory access efficiency. In the dual-channel DDR5 DIMM architecture, the left and right side of the DIMM is each served by an independent channel. The memory modules are installed into matching banks. The overall result is improved concurrency and essentially doubling available memory channels in the system.


Voltage regulation is moved from the motherboard to the individual DIMM, leaving DIMMs responsible for their own voltage regulation needs. This means that DIMMs will now include an integrated voltage regulator.

WHO DOES THIS BENEFIT?

Data centers have the greatest need for the latest memory technology because they must meet the constant demand for lower power requirements, higher density for more memory storage, and faster transfer speeds. With DDR5, servers work more efficiently, and essentially squeeze more ROI out of the investment made in the server.

The appeal of DDR5 will be most notable among those who focus on high-performance systems, both in the deployment of new servers and in replacement of existing infrastructure. Developers of cloud, enterprise, and artificial intelligence applications will benefit from next-generation DIMMs, as will those involved in automotive, networking, and industrial use cases.

While servers will drive initial demand for next-generation DDR, consumers will benefit, too, as in time the technology will make its way into PCs and laptops.



A GAME CHANGER

As noted, the key features driving memory technology are memory density, speed, and faster access. Next-generation, high-performance DDR5 also pushes the envelope in key areas like power consumption and bandwidth to bring new levels of performance to server computing applications. Compared to DDR4, DDR5 delivers 60% faster data transfer speeds. It supports a transfer rate of 41.6 GB/s, equivalent to 11 HD resolution movies, which could be processed in just one second.

DDR5 doubles the bandwidth of DDR4, which tops out at 3.2 Gbps.

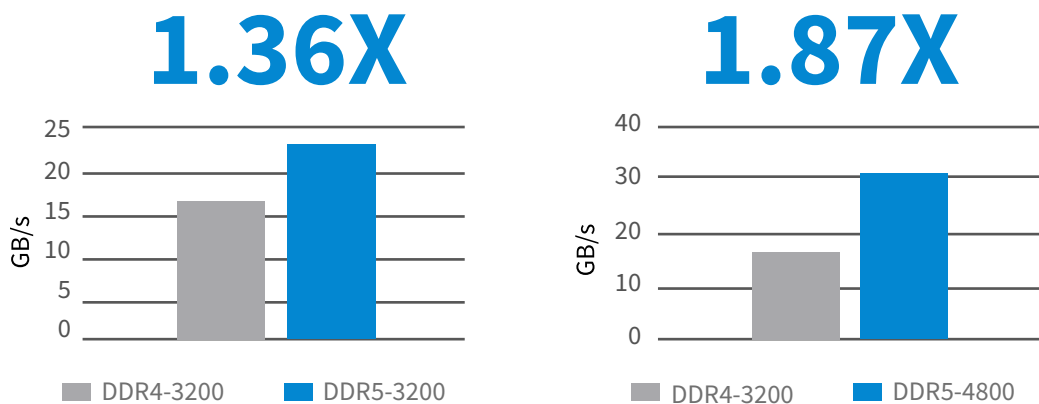


Figure 3: Comparison: DDR4 and DDR5. (Source: Micron)

The standard is architected to enable scaling memory performance without degrading channel efficiency at higher speeds, which has been achieved by doubling the burst-length—which refers to the amount of data read/written after a read/write command is presented to the DDR controller—to BL16 and the bank count to 32 from 16. BL16 allows a single burst to access 64 bytes of data, which is the typical CPU cache line size. It can do this using only one of the two independent channels.

Burst length for SDRAM/DDR/DDR2/DDR3 is 1/2/4/8, respectively. Increasing the burst length reduces the I/O needed to fulfill the same amount of data for a given system access size. This also

results in improved scheduling from the memory controller.

Micron DDR5 DIMMs feature two 40-bit (32 bits plus ECC) independent channels. Having two smaller independent channels improves memory access efficiency. In the dual-channel DDR5 DIMM architecture, the left and right side of the DIMM is each served by an independent channel. The memory modules are installed into matching banks.

DDR5 employs a double data rate Address/Command and Control bus. While the address bus carries the information about the device with which the CPU is communicating, and the data bus carries the actual data being processed, the control bus carries commands from the CPU and returns status signals from the devices. In order to keep the DIMM pin count to 288, which is the case for DDR4, the number of pins dedicated to the Address/Command/Control needed to be reduced.

DDR5 data buffer chips will reduce the effective load on the data bus, enabling the higher-capacity DRAMs on the DIMM without degrading latency.



OTHER NOTABLE FEATURES

There are a lot of architectural features in DDR5 that contribute to taking the performance level up, not just from a data rate perspective but the efficiency with which that is done.

These include:

- Higher reliability with on die ECC and data cyclic redundancy check (CRC) for read and write. On-die ECC reduces the system error correction burden by performing correction during READ commands prior to outputting the data from the DDR5 device. For every 128 bits of data, DDR5 DRAMs will have 8 bits of storage for ECC. Hence, on-die ECC becomes a powerful feature to protect the memory array against single-bit errors. Furthermore, DDR5 introduces an error check and scrub feature where the DRAM device will read internal data and write back corrected data if an error occurred.
- Like its predecessor, DDR5 uses the DDR PHY Interface (DFI), a standard interface between memory controller and PHY, to reduce the integration cost and increase performance and data throughput efficiency. DFI defines signals, timing, and functionality required for efficient communication across the interface. DFI 5.0 has new interfaces to increase speed at which data is sampled, and support for multiple frequency sets resulting in lossless communication to DRAM. It has also added a message interface to improve communication between the memory controller and PHY. In DFI 5.0, the training mode has been transformed to be a PHY-independent training mode; the PHY trains the memory interface without involving the controller.

New and improved training modes help DIMMs and controllers compensate for minute timing differences along the memory bus. This includes a new read preamble training mode, command and address training mode, chip select training mode, and a write leveling training mode that provides the same capability as DDR4, allowing the system to compensate for timing differences on a module.

Memory training occurs on power up, and it is the process whereby the system initializes all the memory installed in a system, does a few rapid tests, organizes it, and then makes it available for use.

- Decision Feedback Equalization (DFE) was added to improve IO speed scalability and enable high-frequency interfaces. In communications, equalizers are used to reduce inter-symbol interference (ISI) to allow recovery of the transmitted symbols. ISI is the effect that a given symbol has on the response from subsequent symbols observed at the receiver. DFE is a filter that uses feedback of detected symbols to produce an estimate of the channel output. This is the first time it's been in the DRAM itself.
- Memory buses such as DDR5 use on-die termination (ODT) modes, eliminating the need for external termination resistors and, as a result, improving signal integrity. It does, however, present a challenge for test probing technology. ODT enables the DRAM to switch between high and low termination impedance. When the termination impedance goes high, the probe impedance needs to be high enough to reduce probe loading.



TECHNOLOGY ENABLEMENT PROGRAM (TEP)

Micron Technology has a comprehensive DDR5 memory enablement program that provides early access to technical resources, products, and ecosystem partners. Called the Micron Technical Enablement Program (TEP), it facilitates the design and development of next-generation computing platforms.

Who's eligible? CPU and ASIC designers, system architects, server manufacturers, OEMs, data center customers, system integrators, and other enterprise computing early adopters that are bringing a related product to market or evaluating a DDR-enabled platform.

By enrolling in the program, approved partners can receive:

- Access to technical resources to aid in product development including data sheets and electrical and thermal models. In addition, partners can find training materials, technical support, technical marketing briefs, and blogs — all located on a single website!

- Access to select DDR5 component and module samples.
- Connection with other ecosystem partners that can aid in the design and development of DDR5 enabled platforms.

TEP reaches beyond Micron products and resources. The program also provides access to other ecosystem partners who can aid in chip- and system-level design. For more on the ecosystem partners enabling DDR5, check out the ecosystem portal. Companies joining Micron in the DDR5 Technology Enablement Program include Cadence, Montage, Rambus, Renesas, and Synopsys.

If you're interested, sign up for the [Micron Technical Enablement program at the **Avent | Micron DDR5 Landing Page**](#). Select the [Request Information button](#), and follow the directions.

MICRON AND AVNET

Micron Technology is the only company that manufactures today's major memory and storage technologies: DRAM, NAND, NOR, and 3D XPoint memory. For nearly 50 years, Micron has partnered with Avnet, a global electronic components distributor with extensive design, product, marketing, and supply chain expertise. Avnet facilitates customer collaboration, support, and quality at every stage of the product lifecycle—whether a product is still in the conceptual stage, already in production, or at any point in between. Offering over 5 million parts online Avnet's position at the heart of the technology supply chain allows it to design, make, supply, and deliver for customers of every size in every corner of the world.

[LEARN MORE >](#)

As a global leader in memory technology, Micron brings more than 40 years of expertise in design, manufacturing, delivery, and support to its global list of partners and customers. Micron's portfolio of DDR5 SDRAM products sets the bar higher than ever before.

[LEARN MORE >](#)

In conclusion, the speed that DDR5 offers is 16x faster than the first ever SDRAM. Compared to DDR4, the DDR5 spec delivers twice the performance and improved power efficiency. Given these significant advantages over DDR4, DDR5 SDRAM has emerged as the near-perfect solution to fill the needs of current and future architectures, meeting the ever-growing demand from data center and cloud environments, as well as HPC and artificial intelligence applications.