Z2DATA

What the AI-Driven Memory Chip Shortage Means for Buyers

How AI demand is impacting specific memory parts, availability, and pricing.

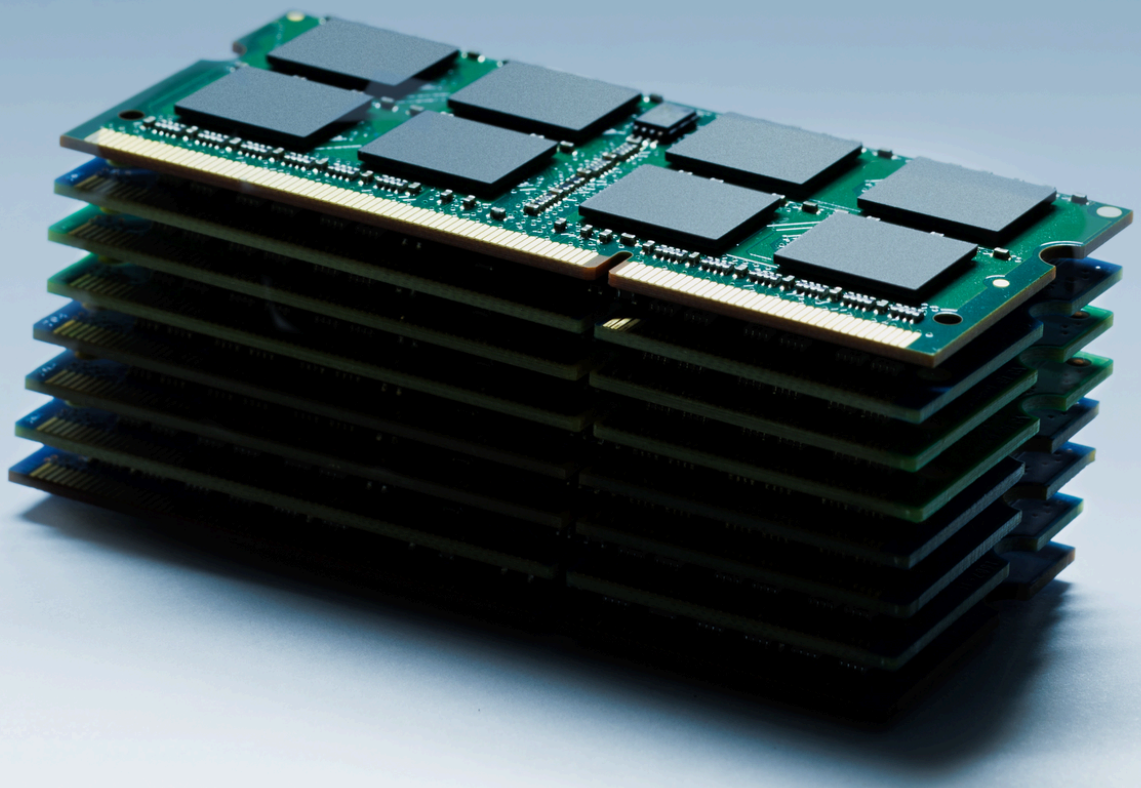

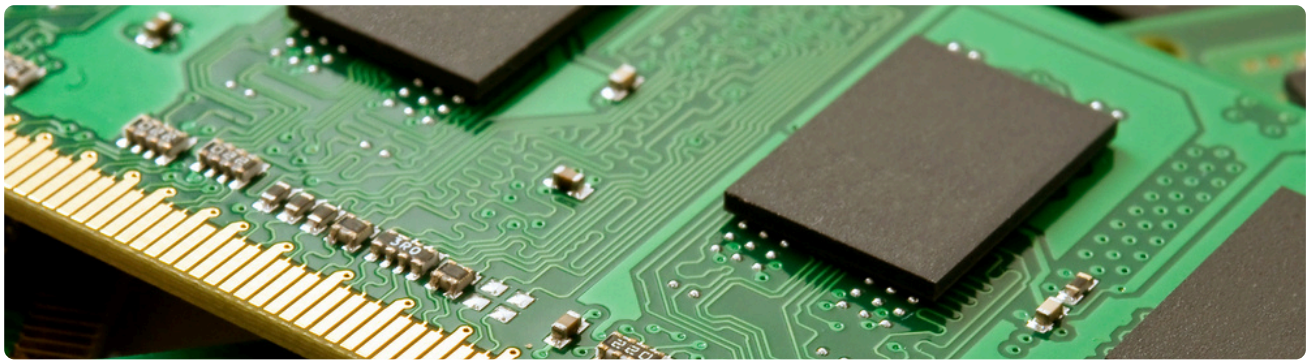

What the AI-Driven Memory Chip Shortage Means for Buyers

It arrived at first gradually, and then all at once.

Prices for memory commodities like dynamic random access memory (DRAM) and NAND started to steadily increase in late 2024 and into 2025, as organizations in industries like automotive, consumer electronics, and cloud computing began observing their RAM costs tick upward, again and again, from one month to the next. Then, toward the end of the summer, those relatively modest price increases accelerated, as costs surged from September onward. All told, media outlets were reporting that by the end of the third quarter, DRAM prices had increased a staggering 172% year-over-year.

> **All told, media outlets were reporting** that by the end of the third quarter, DRAM prices had increased a staggering 172% year-over-year.

In the final months of 2025, the situation grew even more dire. Major technology companies, including “hyperscalers” like Amazon Web Services, Microsoft Azure, and Google Cloud, were absorbing price increases as high as 50% on DRAM orders in October and still only receiving partial fulfillments of their original orders from memory chip manufacturers. In the most extreme cases, contract prices for specific memory commodities like DDR5 were spiking as much as 100% from one month to the next. It will surprise few people to hear that AI data centers and the complex infrastructure they require were among the chief culprits driving these price spikes.

But price spikes are only part of the picture. In this Z2Data analysis, we examine the forces behind these increases, including strategic shifts by major memory manufacturers, accelerated end-of-life (EOL) activity for legacy memory chips, and changes in the geopolitical landscape.

Timeline of Key 2025-2026 Events in Memory Market

- **September 10:** Citigroup issues a warning that the need for DRAM and NAND memory is projected to increase sharply for the fourth quarter of 2025 and into 2026, as memory manufacturers struggle to keep up with exponential growth in demand. Citigroup analysts project a potential supply shortage.

- **September 17:** Memory prices start creeping upward. DigiTimes publishes data showing contract prices for DRAM and NAND increasing up to 20% for the fourth quarter. The publication attributes the price hikes to the downstream effects of production capacity being diverted to AI infrastructure.

- **October 13:** Citing industry data from TrendForce, Reuters reports that memory prices for a variety of DRAM chips increased over 170% in the third quarter year-over-year.

- **November 14:** Semiconductor Manufacturing International Corporation (SMIC), China's largest chip foundry, says fears of memory shortages are prompting customers to delay 2026 orders for other chips, as they worry constrained memory supply could limit their own production. "People don't dare place too many orders for the first quarter next year," SMIC co-CEO Zhao Haijun said, citing uncertainty over how many products memory availability will ultimately support.

- **December 1:** The GM of TeamGroup, a major supplier of memory modules, warns that a supply shortage for memory chips is imminent, and expects availability to tighten even further in the first half of 2026. The GM said he anticipates the shortage lasts into 2027.

- **December 4:** Micron announces it is discontinuing its Crucial line of memory products.

- **January 6:** Technology research firm Counterpoint declares a "hyper-bull" phase in the memory market. The firm cites price increases of up to 50% on memory products in the fourth quarter of 2025, while forecasting prices to surge another 40%-50% in the first quarter of 2026.

- **January 6:** In an interview with Bloomberg, Samsung president and head of global marketing Wonjin Lee said the ongoing memory chip shortage is likely to drive price increases across the electronics industry. "In 2026, there's going to be issues around semiconductor supplies, and it's going to affect everyone, not just Samsung," Lee said, adding that supply constraints are becoming an industry-wide reality.

The Appetite for AI Is Triggering a Scarcity in the Memory Market

To understand how the vast technological ecosystem supporting and advancing artificial intelligence is reverberating through the memory chip market, you need to start with high-bandwidth memory (HBM). HBM is an essential component in AI data infrastructure—and AI accelerators in particular—delivering the high data transfer rates required for AI models to process the enormous datasets that enable their training and inference capabilities.

HBM Demand Is Reshaping Memory Manufacturers' Priorities

Because of HBM's essential utility in the scaling of artificial intelligence infrastructure, AI companies and the data centers they're building all over the U.S. are driving robust, seemingly inexhaustible demand for the memory chips. This has compelled the world's foremost memory manufacturers—including Micron, Samsung, and SK Hynix—to divert their production away from DRAM, NAND, and other more traditional memory commodities, and toward making more HBM. The reallocation of manufacturing resources is gradually reducing the supply and access of the more traditional memory commodities.

> **This has compelled major memory manufacturers**—including Micron, Samsung, and SK Hynix—to divert their production away from DRAM, NAND, and other more traditional memory commodities, and toward making more HBM.

Micron, SK Hynix, and Samsung: A Timeline of Decisions

Dating back to 2024, industry reports and company documentation from the world's largest memory manufacturers showed clear signs that chipmakers were shifting their production capacity to focus on high-bandwidth memory. In July 2024, Chinese financial publication The Commercial Times published a report asserting that the total HBM production for the world's top three memory manufacturers—Micron, SK Hynix, and Samsung—was expected to double in 2025. The report projected that unit production would jump by over 100% the following year, reaching a total capacity of 540,000 units by the end of 2025.

In the spring of 2025, more specific signals were appearing that the world's preeminent memory manufacturers were beginning to prioritize high-bandwidth memory in an even more substantive way. In April, SK Hynix reportedly completed a large-scale conversion of one of its major fabrication plants to ramp up production of HBM packaging lines. The chipmaker carried out the repurposing on M10, one of the company's three main fabrication sites located in and around its headquarters in Icheon, South Korea. According to industry sources, the conversion would allow SK Hynix to expand HBM production from 120,000

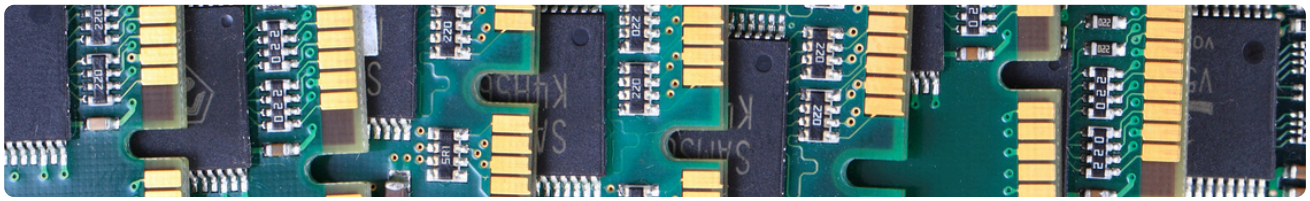

Micron, SK Hynix, and Samsung: A Timeline of Decisions, Cont.

units per month to 130,000. Additional modifications were expected to take place later in the year that would bolster the firm's HBM manufacturing capacity even further.

Finally, in December reports began to emerge that Samsung was seeking to accelerate the timeline for its P4 fabrication site, in Pyeongtaek, South Korea. While that fab remains a multifaceted construction project divided into a number of phases, the company's aim is to use Phase 4 to establish large-scale production of key components utilized in high-bandwidth memory (namely Samsung's 10-nm sixth generation 1c DRAM). While production for Phase 4 was originally scheduled to be completed by 2027, Samsung is now driving toward a completion date of late 2026. There's a clear source of motivation for this production expansion. According to DigiTimes Asia, "Samsung's HBM orders for 2026 are projected to exceed 2025 by more than threefold."

> **While it may not always be immediately evident** from news developments leaked by inside sources, much of the capacity expansion dedicated to HBM production is coming at the expense of traditional DRAM.

HBM Production Comes at the Cost of Traditional DRAM

While it may not always be immediately evident from individual press releases or news developments leaked by inside sources, much of the capacity expansion dedicated to HBM production is coming at the expense of traditional DRAM. In some cases, this takes a highly literal form, in which chipmakers swap out DRAM production lines to make room for HBM manufacturing. But there are subtler ways in which HBM is supplanting DRAM, too.

The more commitments Micron, Samsung, and SK Hynix make to expanding HBM production, the less capable they'll be for meeting the steady and even rising demand for traditional memory commodities like DRAM and NAND. The flipside to every announcement around the expansion of high-bandwidth memory production, in other words, is a further throttling of the memory commodity market, as growing demands are met with stagnant and even shrinking supplies.

DDR and HBM Production Locations: Micron, Samsung, SK Hynix

Memory Manufacturer	Memory Type	Facility Type	Country	Fab/IC Assembly Site
Micron Technology	DDR5/HBM	Fabrication	Japan	Hiroshima
Micron Technology	DDR5	Fabrication	Taiwan	Taichung
Micron Technology	DDR5	IC Assembly	China	Xi'an
Micron Technology	DDR5/HBM (Planned)	Fabrication	U.S.A.	Boise, Idaho
Samsung Electronics	DDR5/HBM	Fabrication	South Korea	Pyeongtaek (P1L, P2L, P3L)
Samsung Electronics	DDR5	Fabrication	South Korea	Hwaseong
Samsung Electronics	DDR5	IC Assembly	South Korea	Onyang
Samsung Electronics	HBM	Fabrication	South Korea	Hwaseong (Line 13, 15, 17)
Samsung Electronics	HBM	IC Assembly	South Korea	Cheonan, Onyang
SK Hynix	DDR5/HBM	Fabrication	South Korea	Icheon (M16, M14)
SK Hynix	DDR5	Fabrication	China	Wuxi
SK Hynix	HBM	IC Assembly	South Korea	Icheon (M10F)
SK Hynix	HBM	Fabrication	South Korea	Cheongju (M15X)

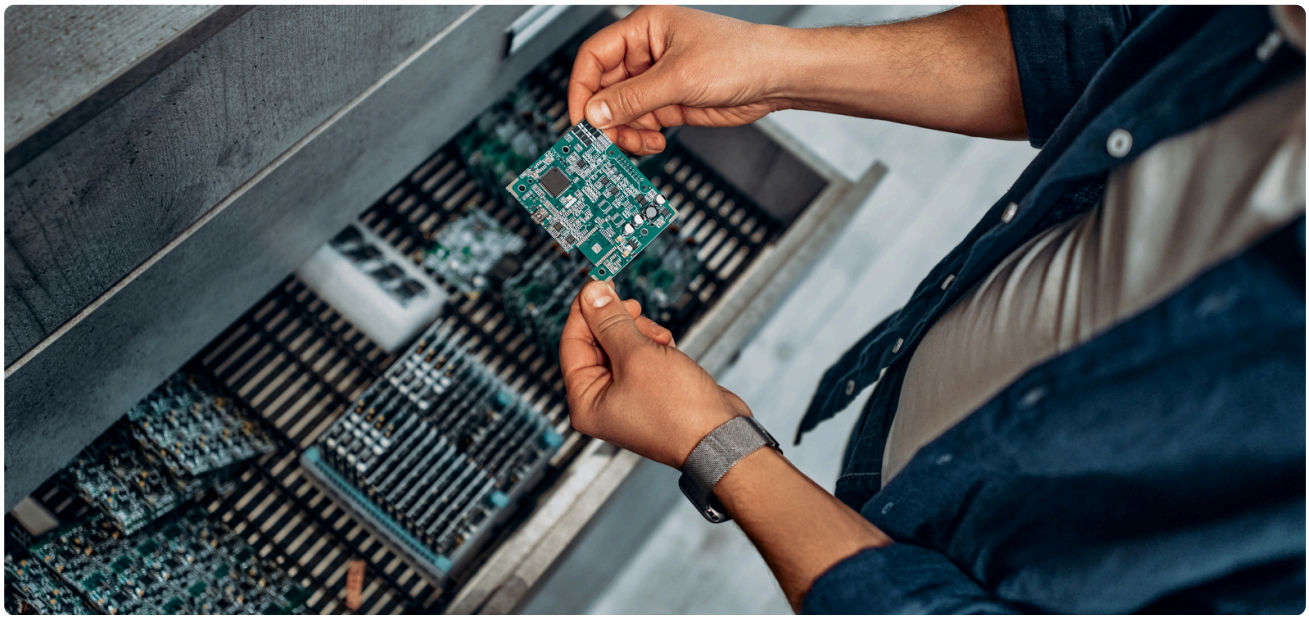

HBM Manufacturing at the Expense of Traditional DRAM, Cont.

As the table on the page before demonstrates, Micron, Samsung, and SK Hynix are now allocating a significant and growing amount of their fabrication and IC assembly capacity to high-bandwidth memory.

> **The rub, of course, is that demand for DRAM, RAM, and NAND** has not diminished for all the other trillion-dollar industries that rely on memory chips to manufacture their core products.

The rub, of course, is that demand for DRAM, RAM, and NAND has not diminished for all the other trillion-dollar industries that rely on memory chips to manufacture their core products. Sectors like automotive, electronics, telecommunications, and cloud computing still require large quantities of traditional memory chips. This has created a stark supply-demand mismatch, one in which a steady demand for commodity memory semiconductors is colliding with an increasingly shallow supply pool.

Memory Manufacturers Are Shifting Away from DDR4 for More HBM

Obsolescence trends can function as another meaningful signal of what semiconductor manufacturers are prioritizing, and what they deem to be most essential to their bottom lines. One place to see the shifting priorities of the world's leading memory manufacturers is in the end-of-life (EOL) trends for traditional memory products like DDR4.

EOL Trends for DRAM: 2023-2025

Technology	2023 Active MPNs	2023 EOL MPNs	2024 Active MPNs	2024 EOL MPNs	2025 Active MPNs	2025 EOL MPNs
SDRAM: DDR4	74%	26%	66%	34%	54%	46%

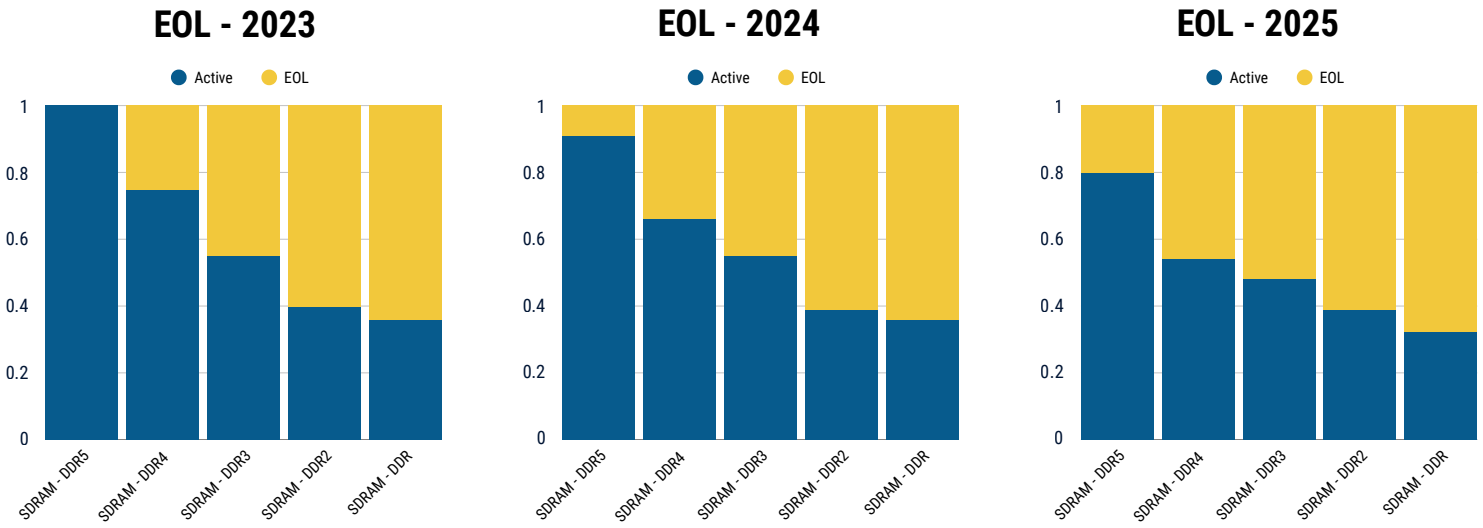

Officially released in 2014, DDR4 chips remained in relatively widespread use in 2023. During that year, nearly three-quarters of the manufacturing part numbers (MPNs) in our database were in active use, with only 26% obsolete. One year later, in 2024, the rate of obsolescence started ticking upward, and over a third of DDR4 chips went into EOL. And in 2025, the number of active DDR4 parts declined even further, with nearly half of these components becoming obsolete.

One year later, in 2024, the rate of obsolescence started ticking upward, and over a third of DDR4 chips went into EOL. And in 2025, the number of active DDR4 parts declined even further, with nearly half becoming obsolete.

Memory Manufacturers Are Shifting Away from DDR4 for More HBM, Cont.

This coincides with the rise of high-bandwidth memory as an in-demand product. It's fair to conjecture that, beginning in 2024, memory manufacturers started making difficult decisions regarding how to allocate their finite production capacity. It's highly likely that firms like Micron, Samsung, and SK Hynix started pulling some of their DDR4 semiconductors out of production in order to expand capacity for the HBM chips that are being deployed for AI data centers all over the country—and, indeed, the world.

> **It's highly likely that firms like Micron, Samsung, and SK Hynix** started pulling some of their DDR4 semiconductors out of production in order to expand capacity for the HBM chips.

This trend is most clearly evident in 2025. The EOL rate for DDR4 jumped an additional 12% this year, and now nearly half of all DDR4 MPNs have an official EOL date. The trendlines are similar for DDR3. While the technology was still crucial to chipmakers in 2023, with nearly 60% of DDR3 MPNs in active use, that figure declined in both 2024 and 2025. Now, less than half of the DDR3 chips in the Z2Data database are still in active use.

Micron Offers a Clear Example of the Industry Shift

Micron's recent decision to discontinue its Crucial line of chips is highly emblematic of the industry's larger shift toward obsoleting traditional memory products. In a press release from December 3, Micron's EVP and Chief Business Officer Sumit Sadana explained the thinking behind the chip manufacturer's pivot. "The AI-driven growth in the data center has led to a surge in demand for memory and storage," she said. "Micron has made the difficult decision to exit the Crucial consumer business in order to improve supply and support for our larger, strategic customers in faster-growing segments." Micron is communicating, in no uncertain terms, its decision to prioritize AI customers and their voracious demand for high-bandwidth memory chips over its older lines of memory products.

> **Micron is communicating, in no uncertain terms,** its decision to prioritize AI customers and their voracious demand for high-bandwidth memory chips over its older lines of memory products.

Signs of Memory Price Shocks Started in 2024

The prevailing media narratives around the memory shortage are not centered on obsolescence trends or strategic shifts to existing production capacity, of course—they're focused on prices. And while our analysis largely corroborates the media stories around stark price increases for memory chips over the second half of 2025, we also found evidence that prices were quietly climbing in 2024, too.

- **From January 7, 2024, to January 5, 2025, the average price of a memory chip in our database increased 160%.**

- **By the end of November 2025, the average price of a memory chip in our database increased over 200% from January 2024 prices.**

Costs and Lead Times Spiking for Top Memory Chips

Our analysis also found that prices and lead times for memory chips from the top manufacturers—Micron, Samsung, and SK Hynix—have been sharply rising since the start of the summer. Conversely, total inventory across all distributors has dropped sharply.

Top Memory Chip Prices and Lead Times

Date	Average Lead Time (Weeks)	Average Price
6/1/2025	13.5	$11.17
8/10/2025	17 (26% Increase)	$13.58
10/26/2025	20.5 (20% Increase)	$15.10
12/7/2025	21.5 (5% Increase)	$19.05

In the six-month period from early June to early December, the total inventory for memory chips from the world's top manufacturers shrunk by nearly 62%, moving from around 34,000 to just under 13,000. The average price for these chips, meanwhile, jumped by over 70% in the same time period. Finally, lead times have also increased, and are now nearly twice as long as they were at the beginning of the year (when they were 12 weeks, on average, for the top memory chips).

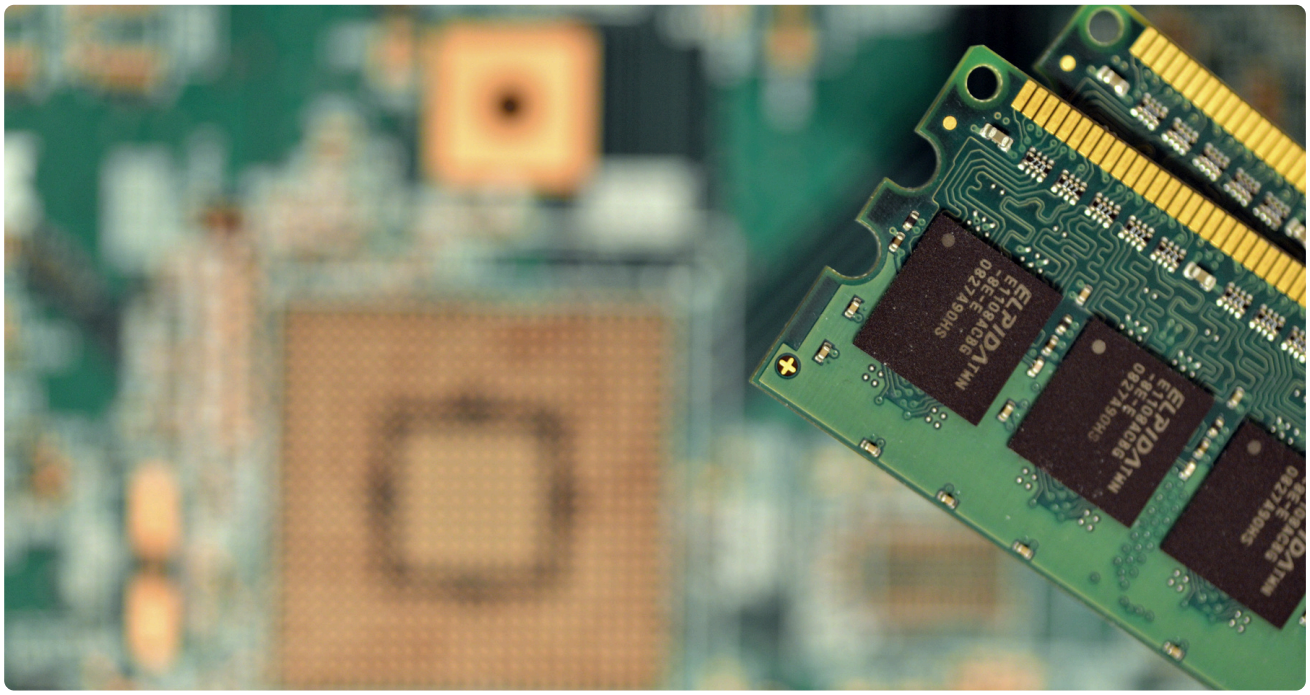

Few Players In the Memory Market Are Alleviating This Chokepoint

It's important to remember, here, too, that the memory manufacturing market is significantly smaller than one might expect. Indeed, it's dominated by just three companies: the aforementioned Micron, Samsung, and SK Hynix. Manufacturing DRAM memory requires highly specialized expertise, and while a larger swath of companies are capable of assembling memory modules, only a few can effectively execute every step in the manufacturing process.

> **As Micron itself explains on its website,** "There are only a handful of semiconductor manufacturers with the capability to produce DRAM chips, and they are Micron (Crucial), Samsung, and Hynix."

As Micron itself explains on its website, "There are only a handful of semiconductor manufacturers with the capability to produce DRAM chips, and they are Micron (Crucial), Samsung, and Hynix." This means that every decision each of these manufacturers makes regarding how to allocate production capacity has significant ramifications for the entire memory market. Simply put, there are no manufacturers waiting in wings to step in and seize on the void left when one of these chipmakers shifts their focus to HBM (although there are smaller memory manufacturers, they do not have the resources or production capacity to scale in a way that effectively compensates for the major players shifting to HBM). Instead, the supply tightens, and prices climb.

Memory Shortages: A Problem Unlikely to Go Away

With so many developments, disruptions, and seemingly abrupt shifts in the semiconductor manufacturing industry, it can be difficult to distinguish between temporary concerns and serious, enduring issues. But it may be useful to compare the emerging memory chip shortage to another recent crisis—the Nexperia ownership dispute.

The ownership dispute over Dutch chipmaker Nexperia drew attention for its geopolitical overtones and briefly disrupted the semiconductor supply chain, particularly in the automotive sector. While some automakers adjusted operations, the worst effects were short-lived, as Nexperia resumed exports from China and continued fulfilling contracts, despite lingering complications.

It's unlikely that the supply and demand issues around memory chips, on the other hand, will have a similarly swift resolution. As we outlined earlier, this is a systemic issue, one in which surging demand for high-bandwidth memory is compelling manufacturers to reallocate fab capacity away from traditional memory. Given the foundational issues at play, it's no surprise that some industry insiders foresee a shortage lasting into 2027.

Taking Steps to Mitigate the Impact

With forecasts suggesting tightness could last into 2027, the key question for manufacturers is no longer whether the memory market will stay volatile, but how to reduce exposure to price spikes, constrained availability, and growing EOL risk. The most effective approach is to shift from reactive firefighting to proactive monitoring and mitigation.

Z2Data's risk management platform offers a number of capabilities that can help businesses navigate the rapidly evolving memory market:

- **Z2Data's Part Risk Manager** tool lets companies to see prices, inventory, and lead times across all distributors in real time, including both authorized and independent sellers.
- **Z2Data lets users set up real-time alerts** at the component level, establishing triggers for variables like lead time, lifecycle changes, and changes to manufacturing sites.
- **Z2Data's Lifecycle Forecasting** gives organizations an accurate, effective estimate of the obsolescence date for their parts, helping them understand the memory chips at the greatest risk of EOL.

To learn more about how Z2Data can help you effectively mitigate this growing supply chain risk, schedule a demo with one of our product experts or contact your representative.

Schedule a Demo