

IMPULSE EMBEDDED

Why a Volatile Market Is Changing Your Industrial System Quotes & How to Stay Ahead

WWW.IMPULSE-EMBEDDED.CO.UK



Memory, storage and CPU prices are moving again.

In recent weeks, you may have found yourself wondering:

- “Why has our system price gone up so much since the last quote?”
- “Why is my quote only valid for a short time?”
- “Why has the lead time doubled?”

From our side, the pattern is very clear. Prices and lead times for DRAM, NAND Flash, SSDs and some CPUs are moving quickly, and in many cases in the wrong direction for project budgets.

This is by no means limited to a single manufacturer. It is the result of a global shift in the semiconductor market, driven by a combination of AI demand, data-centre investment and the production cuts memory makers made in 2023 to cope with oversupply.

The purpose of this article is to explain how we got here, what is happening right now in practical terms, and how we are handling it at Impulse so you can plan your industrial and embedded projects with fewer surprises.

Contents:

- How Did We Get Here?
- What’s Happening Now & Why
- What This Means for Industrial & Embedded Projects
- How We’re Handling It
- A Practical Word on Planning Ahead





How Did We Get Here?

To understand the current volatility, it helps to look at the past few years as three distinct phases:

The Pandemic Shortage

The first phase is familiar to anyone who tried to source components around 2020 - 2022. A combination of pandemic disruption, strong demand for electronics and earlier under-investment in capacity led to broad semiconductor shortages. Lead times for many devices, including microcontrollers and memory, stretched from a normal 8 - 12 weeks to 40 weeks or more in some cases. That period triggered heavy investment in new capacity, but it also left the industry cautious about overshooting again.

The Deep Downturn in 2023

The second phase arrived faster than many expected. By 2023, demand for PCs and smartphones had cooled, and the big memory suppliers found themselves with too much stock and falling prices. Gartner reports that memory revenue declined by 37% in 2023, with DRAM revenue down 38.5% and NAND Flash revenue down 37.5% year on year.

When revenue falls that far, suppliers act. Companies such as Micron, Samsung and SK Hynix cut production of both DRAM and NAND and reduced capital expenditure on new equipment and fabs. One analysis notes that wafer starts for DRAM and NAND were reduced by about 30% at Micron, with those reductions extending well into 2024.

In simple terms, the industry took capacity out of the system to stop prices collapsing further. That decision was rational in 2023. It also set the stage for the current tightness once demand returned.

AI and Data Centres on a Lean Supply Base

The third phase is where we are now. AI workloads moved from proof-of-concept into large-scale deployment just as the memory industry had cut wafer starts and worked down inventories after the 2023 downturn. Cloud providers and hyperscalers started ordering huge volumes of DRAM, high-bandwidth memory (HBM) and enterprise SSDs to build AI servers and modernise data-centre infrastructure.

Suppliers have responded by shifting more of their output into higher-margin products such as HBM and DDR5 for AI and servers, and by gradually phasing down older DDR4 and legacy NAND lines as a share of total production. At the same time, memory inventories are far lower than they were during the 2023 glut, leaving much less buffer in the system when demand spikes.

That combination of stronger AI and data-centre demand, tighter production capacity and much lower stock levels is the backdrop to the price rises and lead-time changes industrial customers are seeing today. The next section looks at what is happening now in more detail.



What's Happening Now & Why

In practical terms, that backdrop now shows up as rapidly rising contract prices, tighter allocation on certain product lines, and a stronger focus on high-margin AI and data-centre customers.

DRAM and DDR4: Squeezed Between AI and End of Life

Contract prices for DRAM and NAND were initially estimated by market trackers to be up by around 15 - 20% in the fourth quarter of 2025 compared with the previous quarter, with some supplier and analyst reports pointing to increases of up to 30% for particular configurations. TrendForce's own data for conventional DRAM in 4Q25 points to price increases roughly in the 18 - 23% range for many categories as AI-related demand has accelerated.

The result, as we've seen it, has been even higher than these estimations, and a recent report signals that memory prices have, and will, exceed prior expectations. Samsung themselves have reportedly hiked the prices of their server chips by 30% - 60% compared to September 2025.

At the same time, suppliers are converting lines from "conventional" DRAM to HBM and newer DDR5 products. For industrial customers that still rely on DDR4, that creates a double squeeze. On one side, total DDR4 output is being reduced; on the other, the remaining output is often prioritised for large, long-term customers and placed on allocation. TrendForce's forecast that DDR4 capacity will drop to around one fifth of 2025 levels by 2026 illustrates how quickly this shift is happening. Simultaneously, DDR5 is being heavily prioritised for AI and hyperscale projects, leaving little supply for everyone else.

From a project perspective, that is why we are seeing:

- Fewer options on certain DDR4 densities and speed grades.
- Even less supply available for DDR5.
- Longer and more variable lead times.
- A stronger push towards newer memory platforms in some CPU and embedded board families.

NAND Flash and SSDs: AI Storage is Soaking Up Wafer Supply

On the NAND side, the numbers are even more striking. TrendForce reports that in November 2025, contract prices for NAND wafers rose by between 20% and more than 60% month-on-month, depending on the density and type, as suppliers focused their capacity on high-margin enterprise and premium products.

That shift is closely linked to AI storage demand. Enterprise SSD volumes for AI workloads are rising quickly, and AI server SSD procurement is forecast to grow at an average annual rate of over 60% for the next few years. As hyperscalers sign large contracts for high-capacity drives, the wafers that underpin more "ordinary" client and industrial SSDs become scarcer and more expensive.

Even if you are buying a compact industrial PC with a relatively modest SSD, you are competing for capacity with data-centre customers ordering racks of AI servers. That is why SSD prices and lead times are now moving more in step with data-centre cycles than with PC cycles.



CPUs, SoCs and Embedded Platforms

Processors are not immune to these trends, although the picture is more mixed. Leading-edge CPU and SoC production shares capacity with AI accelerators at advanced foundries. As more of that capacity is booked for GPU and accelerator devices, room for other products tightens and pricing becomes less flexible. At the same time, mature-node foundries are in demand from automotive, industrial and IoT customers that still use older geometries.

For industrial and embedded platforms this often shows up not as a headline “shortage of CPUs” but as:

- Longer quoted lead times on newer embedded boards that use the same nodes as popular consumer or server CPUs.
- Reduced flexibility on last-minute changes to CPU or chipset selections.
- A more frequent need to talk to us about long-term availability when you are planning a design intended to live for many years.

How Long Could This Last?

Forecasting is always uncertain, but several signals point in the same direction. [Reuters recently described the situation as a new global memory supply chain crisis,](#) and analysts expect memory tightness to continue through 2026, and some of the new memory fabs being announced today are not expected to reach meaningful volume until [2027 or 2028.](#)


From our point of view at Impulse, the sensible working assumption is that volatility in DRAM, NAND and some processor families will remain elevated for at least the next twelve to twenty-four months, with the possibility of further spikes along the way.



What This Means for Industrial & Embedded Projects

The most immediate impact for our customers is that quotes for systems and components containing DRAM, Flash, SSDs and certain CPUs are higher and less stable than they were even six months ago. Prices that used to hold for sixty or ninety days may now only be valid for a much shorter period on some parts because upstream contract prices are moving more quickly. Lead times that sat comfortably in the twelve to sixteen week range can extend beyond twenty-six weeks where parts move onto allocation or where a manufacturer has shifted capacity away from a particular line.

You may also notice more conversations about alternatives. That might mean proposing a slightly higher density DRAM device where the original density is constrained or suggesting a different SSD series that uses a more available NAND configuration. None of those recommendations are made lightly. The aim is to keep industrial projects viable in an environment where the “easy” choice on paper may simply not be available in the quantities or timeframes you need.



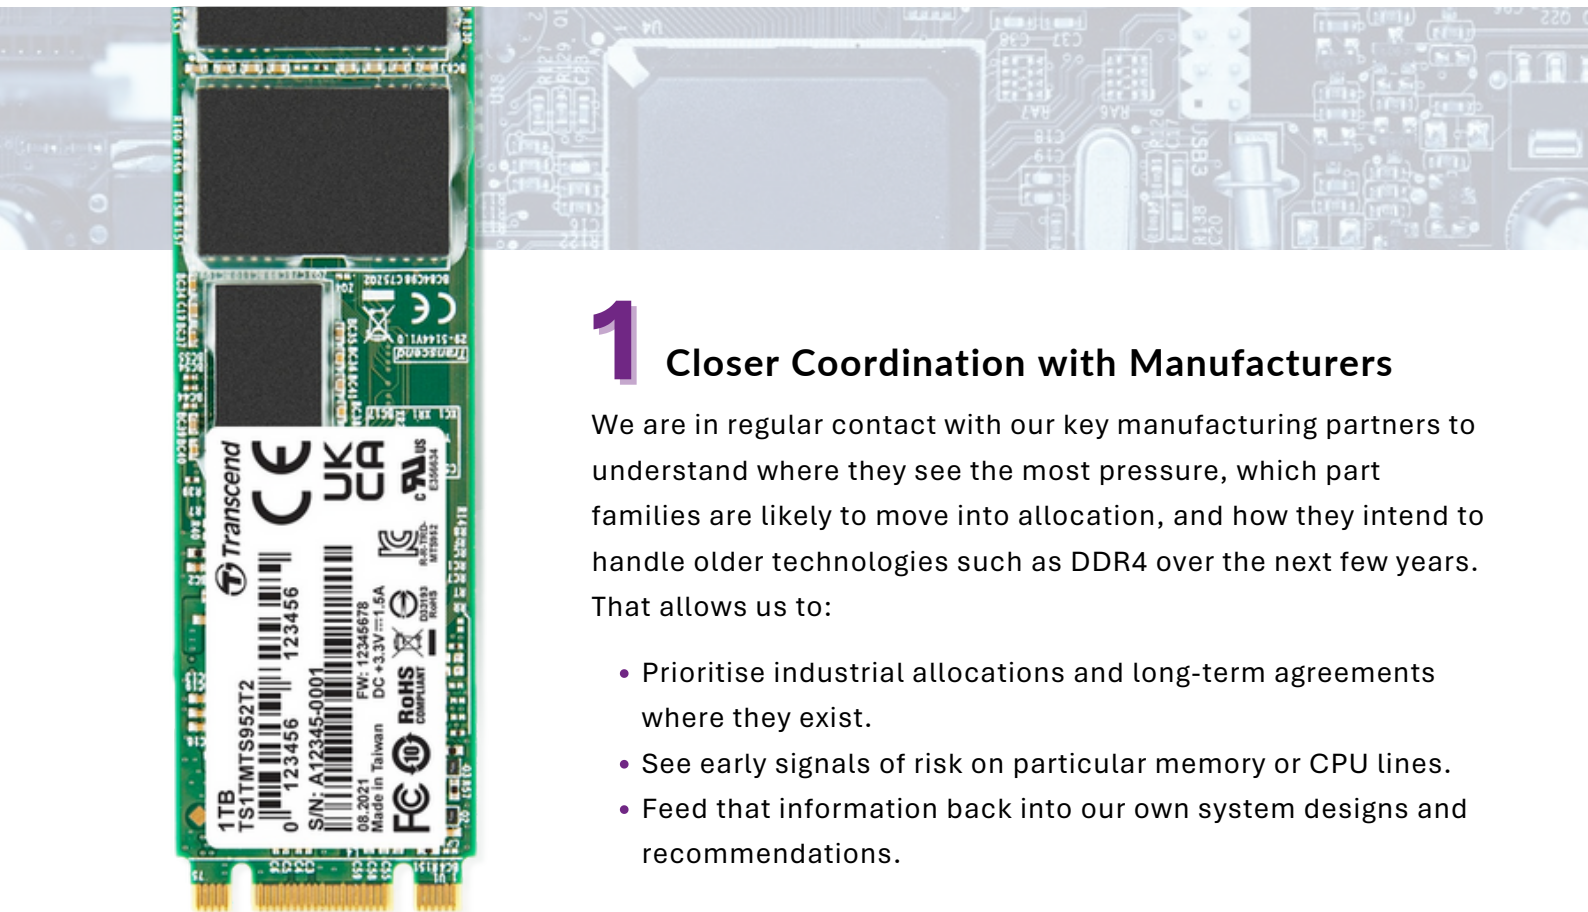
It's worth stressing that not every product is affected to the same degree. Long-lifecycle industrial platforms which already sit under firm supply agreements, and some lower volume or niche configurations, have been more stable so far. However, the direction of travel across the broader market is clear enough that we now treat memory and CPU selection as a supply-chain decision as much as a technical one.



How We're Handling It

Our role in situations like this is to act as a bridge between global semiconductor cycles and the day-to-day needs of UK OEMs, system integrators and operators.

There are four main strands to how we are approaching the current volatility.



1 Closer Coordination with Manufacturers


We are in regular contact with our key manufacturing partners to understand where they see the most pressure, which part families are likely to move into allocation, and how they intend to handle older technologies such as DDR4 over the next few years. That allows us to:


- Prioritise industrial allocations and long-term agreements where they exist.
- See early signals of risk on particular memory or CPU lines.
- Feed that information back into our own system designs and recommendations.

2 Ongoing Monitoring of Availability

Internally, we treat DRAM, NAND, SSDs and key processor families as a “watch list”. Market data, pricing changes and lead-time movements are reviewed regularly, and our account managers receive updates on where the risk is greatest.

The aim is not to flood you with market commentary, but to make sure that when we say “this is a good time to place that order” or “this part is moving into a tighter supply position”, it is based on current data rather than guesswork.





3 Design and Selection Support

Because we design, build and support industrial computing solutions, we are used to looking at bills of materials through both a technical and a supply-chain lens.

If you are in the early stages of a project, we can help you:


- Identify memory and storage options that balance performance and supply-chain resilience.
- Evaluate the trade-offs and availability involved in using DDR4 or DDR5, or choosing or between different SSD lines.
- Consider life-cycle and long-term availability alongside pure performance numbers.

In some cases, that may mean recommending newer platforms, such as the Intel Core Ultra-based systems we now offer, where we believe the underlying CPU and memory technology roadmap aligns better with the direction manufacturers are taking.

4 Commercial Planning for Critical Projects

Finally, we are spending more time with customers on commercial structures that help smooth out volatility. That can include scheduled orders, phased deliveries aligned with project milestones, or managed buffer stock for truly critical programmes where a disruption would be very costly.

We will not push a one-size-fits-all model here. For some customers, a simple change in ordering behaviour combined with more regular communication is enough. For others, especially where large fleets or long-term service contracts are involved, a more formalised plan makes sense.





A Practical Word on Planning Ahead

None of this means that industrial and embedded projects need to stop or that every design must be rewritten around the latest memory technology. It does mean that a little more planning goes a long way.

If you have projects on the horizon that depend on DRAM, Flash, SSDs or specific CPUs, the most useful step you can take is to bring us into the conversation early. Even a rough view of likely volumes and timing over the next six to twelve months helps us to reserve capacity, choose more resilient part combinations and flag any higher-risk items before they become a problem.

Equally, if you are sitting on quotes for approved projects that have not yet been converted into orders, it is worth reviewing those with us. In the current environment, waiting until the last possible moment often means paying more or accepting less flexibility on configuration.



Talk To Us

Memory and processor markets are going through another uncomfortable adjustment, and it is entirely reasonable to be concerned about the impact on your projects.

If you are an existing customer and would like to talk through the implications for a current or future project, please get in touch with us. We can review your bill of materials, highlight any components that look exposed, and discuss design or commercial options that make sense for your situation.

The underlying technologies will continue to improve and the supply side will eventually catch up. In the meantime, good information and early planning put you in a much stronger position, and that is where we are focused.

If you would like to get in touch, simply contact us on:

☎ +44(0)1782 337 800

✉ sales@impulse-embedded.co.uk

Or visit our [contact us](#) page.