2026 FUSION WORLDWIDE Quarter 1

State of the Industry

Navigating the 2026 Supercycle

When Everything Goes Short:
A Procurement Leader's Guide

Contents

The 2026 market has entered a supply-driven supercycle. Demand tied to AI infrastructure is expanding faster than the systems required to support it, pushing constraints beyond semiconductors into the materials, components, and infrastructure needed to deploy them.

What began as a silicon and memory conversation has evolved into a broader supply imbalance across the entire enabling stack.

For procurement leaders, the challenge is no longer securing a single category. It is identifying where constraints emerge next and how they cascade through deployment timelines.

What's Short Isn't What You Think

The real bottleneck in 2026 is not the component in the headline. It is the sequence of dependencies required to deploy AI infrastructure tightening at once.

For two years, procurement strategies focused on securing silicon and memory allocation. That framing is now incomplete. Silicon and memory are only two nodes in a supply chain that also spans specialty gases, substrate materials, power components, optical systems, and thermal infrastructure before anything reaches a data center floor.

Every one of these nodes is under pressure, several in active shortage. Unlike prior cycles, where constraints were concentrated in finished semiconductors, this environment is distributed across upstream inputs that many procurement teams are not actively tracking.

Driven by AI demand, hyperscale buildouts are pulling supply forward across memory, power, and infrastructure. Across every layer required to deploy AI systems, the signal is consistent: committed, constrained, or tightening.

Category	Status
TSMC N3 Logic Wafers	Fully committed. AI projected to consume 60% of output in 2026, 86% in 2027.
High-Bandwidth Memory (HBM)	Fully allocated through 2026. Contracts locking in high-teens to low-20% price increases.
Helium	Spot prices up 70 to 100% after Iranian drone strikes took Qatar's Ras Laffan offline. 30% of global supply removed.
T-Glass Substrates	Rationed by a single Japanese manufacturer. No new capacity until late 2026.
ABF Substrate Film	Single commercial-scale producer. Tightening as AI packaging absorbs increasing share of output.
Power Management ICs	40 to 60 weeks.
800G Optical Transceivers	36 to 50 weeks. Pricing described only as severe.
Ribbon Fiber	60+ weeks. No confirmed ceiling.
Liquid Cooling CDUs	Production-constrained across every major vendor.

"The shortage is not concentrated in one category. It runs through the entire stack. Across the procurement organizations we support, sourcing pressure is appearing simultaneously in power, optics, substrates, and thermal infrastructure often before customers realize those categories are gating their deployment timelines."

Andrew Czuczwa
Market Research Manager, Fusion Worldwide

$6.7T

Projected Global Data Center CapEx by 2030 (McKinsey, April 2025)

~60%

Share of TSMC's N3 Output Consumed by AI Accelerators in 2026 (SemiAnalysis)

30%

Share of Global Helium Supply Offline After Iranian Drone Strikes (Qatar, March 2026)

~70-100%

Spot Helium Price Increase since Ras Laffan Went Offline (CSBC, March 2036)

"Most of the memory from the top players is going directly to AI infrastructure, but many other markets are starved today because there is no capacity left for them."

Sassine Ghazi, CEO, Synopsys
Q4 2025 Earnings Call, January 2026

The Shortage Has Been Moving Upstream for Three Years. Here's Where It Stands Now

To understand where the shortage is today, you need to understand how it got here. Each year since 2024, the binding constraint has moved one layer deeper into the stack.

The shortage has shifted to the material and infrastructure layer — not finished inventory or a single process step.

It now sits in the specialty gas used to etch wafers, the fiberglass reinforcing substrates, the films forming chip packaging, the passive components supporting power delivery, and the infrastructure enabling deployment.

These inputs have no viable substitutes, cannot be stored beyond limited windows, and in several cases come from a single commercial-scale producer.

What has changed is not just where the shortage sits, but why. Demand is being driven by infrastructure buildout, not consumer recovery. That demand consumes more memory, power, networking, and thermal capacity per deployment than any previous compute cycle.

Constraint Category	2024	2025	2026
Physical Infrastructure	Active	Resolved	Resolved
Advanced Packaging (CoWoS)		Active	Easing
Materials & Infrastructure Inputs		Building	Active

What a System-Wide Shortage Actually Looks Like

Eight categories are in active shortage as of Q1 2026. Each has a distinct cause, confirmed lead time range, and measurable pricing movement.

For most categories, meaningful supply relief does not arrive within the 2026 planning window.

This is not a series of isolated constraints. It is a connected system, where pressure in one category amplifies constraints in another.

Silicon: TSMC N3 Is Fully Committed Through 2026

Every leading-edge AI accelerator shipping in 2026, NVIDIA Rubin, AMD MI350, Google TPUv7, and AWS Trainium3, is built on TSMC's N3 process node. The industry converged on a single node simultaneously, faster than TSMC anticipated. Effective N3 utilization is expected to exceed 100% in the second half of 2026. New cleanroom capacity takes years to build. AI customers receive allocation priority because they carry the highest margins and multi-year purchase commitments. Smartphone customers are being pushed to N2 or extended product cycles. For procurement organizations outside the top tier, N3-based components will remain difficult to secure at any price through 2027.

This is the upstream constraint that makes every downstream shortage worse. When wafer supply is the binding limit, no amount of packaging capacity or substrate material resolves the bottleneck. The distributed shortage documented in this report starts here.

Source: SemiAnalysis Foundry Model, March 2026.

> "We continue to observe strong AI-related demand for our leading-edge technologies."

C.C. Wei, Chief Executive Officer, TSMC
Q4 2024 Earnings Call, January 2025

Memory: HBM Is Sold Out. DRAM Is Tightening Behind It

High-Bandwidth Memory (HBM) has become the second major constraint after wafer supply, and the two are structurally linked. HBM consumes roughly three times more wafer capacity than commodity DRAM per bit, a ratio expected to widen to four times with HBM4. Every wafer start directed toward HBM pulls capacity from commodity DRAM, tightening both markets simultaneously. Samsung, SK Hynix, and Micron are all at full capacity. All 2026 HBM production is allocated. HBM content per accelerator is rising at the same time: NVIDIA Rubin carries 50% more than Blackwell, with Rubin Ultra adding a further 4x. AMD MI400 is up 50% from MI350. Google TPUv8AX and Trainium3 are both moving to 12-Hi HBM stacks. More memory per chip, fewer chips available, all production sold.

Behind HBM, server DRAM is tightening. AI workloads are driving CPU-to-GPU ratios higher, increasing DDR5 demand in AI server deployments. DRAM supplier inventories have fallen to historically low levels. The memory market is not in a single-product shortage. It is in a system-wide tightening with no near-term release valve. This is the second layer of the distributed upstream constraint: not just fewer chips, but less memory per chip that can be produced.

Sources: Micron Q1 FY2026 Earnings Call, December 2025. SK Hynix Q4 2025 Earnings Call, January 2026.

> **"Industry demand is greater than supply for both DRAM and NAND. We expect tightness to persist through and beyond calendar 2026."**
>
> **Sanjay Mehrotra, President & CEO, Micron Technology**
> Q1 FY2026 Earnings Call, December 17, 2025

What We're Watching: Inference Efficiency

Improvements in model efficiency are reducing per-query compute requirements. If this trend accelerates, it could moderate demand growth for HBM and N3 wafers at the margin. This is not yet a supply relief signal. It is a variable worth tracking. Efficiency has historically expanded AI usage rather than reducing total compute consumed.

Helium: Input Procurement Teams Are Not Tracking

In 2025, South Korea sourced nearly 70% of its helium from Qatar. When Iranian drone strikes hit QatarEnergy's Ras Laffan Industrial City in March 2026, the world's largest LNG export facility and a byproduct producer of roughly 30% of global helium supply, the impact was immediate. The facility went offline, and spot prices surged 70 to 100% within days. Fabs in South Korea and Taiwan began drawing down buffer inventories, which industry experts estimate at roughly 45 days before operational constraints emerge.

Helium is critical to EUV lithography, wafer cooling, and cleanroom leak detection, with no viable substitute in advanced semiconductor manufacturing. As supply tightens, fabs reduce production rates to preserve process integrity, directly constraining wafer output. For manufacturers already operating within limited capacity environments, the risk is not theoretical. If disruption extends beyond 60 to 90 days, it introduces an additional layer of throughput pressure on top of existing structural constraints.

While the Strait has since reopened, the situation remains fluid. The exposure is clear, and the market continues to monitor conditions closely.

Sources: Helium Shortage 2026 • CNBC, March 19, 2026. • Resilinc Special Report, March 2026. Tom's Hardware, March 2026. tomshardware.com

T-Glass & ABF: The Materials Holding the Chip Together

AI chip substrates depend on two raw materials most procurement teams have never tracked. Both are in active shortage. Both are produced at commercial scale by a single company each.

T-Glass, a specialized fiberglass produced primarily by Nitto Boseki, prevents substrate warping under the thermal load of high-power AI chips. Standard materials fail this test. Nittobo issued a 20% price increase in August 2025 with no new capacity until late 2026, contributing to a projected 10 to 20% supply gap across Taiwan's substrate manufacturers.

ABF, Ajinomoto Build-up Film, forms the insulating layers in advanced chip substrates. As AI packaging absorbs an increasing share of Ajinomoto's output, availability for other applications tightens. A 50% capacity expansion has been outlined, but it does not address demand that is constrained today.

The bottleneck is not at the finished component level. It is one layer upstream, in the raw materials required to produce it — where most procurement strategies have no visibility.

Sources: Astute Group, February 2026. astutegroup.com • Goldman Sachs Research, T-Glass Shortage Analysis, 2026 • DigiTimes, November 2025.

Power Management ICs: Demand Shifted From Auto to AI. Supply Did Not

AI servers are drawing 10 to 15 times more power than traditional server architectures, fundamentally changing what is required to support them.

Three years ago, a 400-watt server chip was considered high-power. Today, AI accelerators routinely exceed 1,000 watts of Thermal Design Power. The power management ICs and passive capacitors required to handle that load were not in volume production at the scale now needed when this buildout began. The same fab capacity now producing AI-optimized VRMs and power ICs is the capacity that previously supplied automotive MOSFETs, industrial switching regulators, and consumer power components.

Across the procurement intelligence we track, the displacement effect is visible in lead time data. Automotive eMMC lead times have stretched into 2027. Automotive MCU lead times remain at 20 to 45 weeks. Supply that used to serve those markets has been reallocated, and those markets have no near-term recovery path. The pricing signal confirms it: Infineon is implementing broad increases effective April 1, 2026, forecasting a 67% jump in AI data center power revenue year over year. Texas Instruments and ADI have implemented 10 to 30% increases across thousands of SKUs. AI-grade MLCC channel inventory has dropped below one month of supply, with 15 to 20% spot price increases confirmed in Q1 2026.

Networking & Optics: AI Cluster, From Network to Compute Platform

Every connection in an AI cluster runs through switching fabric and optical interconnects that are themselves in shortage. The shift to 800G and 1.6T Ethernet has pushed transceiver lead times to 36 to 50 weeks, with networking ASICs and switches at 40 to 56 weeks, up from a 12 to 16 week baseline. A modern GB300 NVL576 mega-cluster requires 640+ switches, compared to 12 for a 2022 HPC cluster, while ribbon fiber has crossed 60 weeks as hyperscale demand continues to pull supply forward.

At the component level, shortages in isolators and EMLs are creating production bottlenecks, while supplier allocation is concentrating among top-tier customers following recent multi-billion-dollar investments, limiting availability across the broader market.

Networking is where the constraint becomes visible in deployment timelines. Secured compute. Secured memory. Deployment still delayed, because the fabric connecting those systems sits on a 40 to 56-week queue. The constraint is not isolated. It is systemic.

Liquid Cooling: This Is Not Optional Equipment

The GB300 NVL72 rack draws 132 to 142 kilowatts. Air cooling at that density is not an engineering tradeoff. It is a physical impossibility. Direct-to-chip liquid cooling is not optional equipment. It is the condition under which current-generation AI infrastructure can operate at all. CDU production is now a gating constraint on par with silicon and memory: organizations that have secured compute, power contracts, and data center space still cannot turn their systems on without it. The market is projected to grow from $2.8 billion today to over $21 billion by 2032. That trajectory does not resolve the imbalance. It extends it.

What we hear consistently from procurement leaders: thermal infrastructure is treated as an infrastructure decision, not a supply chain decision. By the time it enters the procurement conversation, the queue is already months long. This is a category where early positioning has an outsized impact on deployment timing.

Power Transformers: The Infrastructure Shortage Nobody Is Tracking

A data center without a power transformer is a concrete shell. A facility cannot receive power, run servers, or operate cooling systems until large-scale power infrastructure is in place to convert high-voltage utility power to the levels facilities can safely operate on. Large power transformer lead times now average 128 weeks, with specialized semiconductor fab units stretching to 200 weeks, according to Wood Mackenzie. The supply deficit for power transformers is projected at 30% for 2025. Prices have increased 77% since 2019 for large units and up to 95% for distribution transformers. Major suppliers are implementing 20 to 30% additional price increases in 2026, and some manufacturers are declining to quote on large-scale projects because they cannot meet the timelines.

Data centers, semiconductor megafabs, renewable energy buildouts, EV charging infrastructure, and utility grid modernization are all competing for the same transformer supply. US power transformer demand surged 116% between 2019 and 2025. Imports now account for an estimated 80% of US supply. The National Infrastructure Advisory Council has cited five-year wait times at individual manufacturers. Treating power transformer procurement as a facilities decision is a planning error with schedule consequences.

128 Weeks
Average Lead Time
Large Power Transformers
(Wood Mackenzie, 2025)

200 Weeks
Lead Time
Specialty Semiconductor
Fab Transformer Units

+77%
Price Increase
Large Power Transformers
since 2019

30%
Projected US Power Transformer Supply Deficit
(Wood Mackenzie, 2025)

This Cycle is Structurally Different, and Old Relief Mechanisms Do Not Apply

In 2021 and 2022, the semiconductor shortage was driven by a surge in consumer electronics demand. Supply chains overbuilt inventory to compensate. When demand normalized, that inventory worked its way back through the market. Lead times compressed. Prices fell. The cycle corrected through inventory drawdown and normalization.

That mechanism does not exist in this cycle. The materials and process inputs that are short today, N3 wafer starts, HBM stacks, helium, T-glass, ABF, are consumed in production. When a wafer is etched, the helium is gone. When a substrate is laminated, the ABF is gone. There is no overhang waiting to release. Supply correction here requires new capacity, and new capacity in these categories takes years to build.

This cycle is structurally different in three ways:

No Phantom Inventory: The materials in shortage, N3 wafer starts, HBM stacks, helium, T-glass, ABF, are consumed in production. There is no overhang waiting to flush through.

Demand is Infrastructure-Driven, Not Consumer-Driven: Hyperscalers are committing capital at multi-year horizons. Google, Microsoft, Amazon, and Meta have collectively signaled hundreds of billions in 2026 data center investment. This is committed capital with committed timelines.

Material Input Constraints: With two-to-four-year expansion cycles, T-glass production cannot be accelerated by adding a shift. Helium supply cannot be expanded by calling a supplier. A new N3 fab cannot be built in a year.

Supply Concentration Makes Every Constraint Systemic

In the 2021 to 2022 cycle, constraints were spread across hundreds of components and dozens of suppliers. In this cycle, each critical input traces back to one or two sources. TSMC produces the majority of leading-edge logic wafers. Samsung, SK Hynix, and Micron collectively supply all HBM. Nitto Boseki produces most of the world's commercially qualified T-glass, and Ajinomoto produces most of its ABF substrate film. Qatar's Ras Laffan supplied roughly a third of global helium before going offline in March 2026. A handful of OSAT providers control nearly all advanced packaging capacity. These are not obscure inputs with easy substitutes. They are foundational materials with no commercial alternatives at scale, and in 2026 every one of them is simultaneously constrained.

That is what makes this cycle systemic. When Nittobo issues a 20% price increase, every substrate manufacturer in the world absorbs it on the same day. When Ras Laffan goes offline, every advanced fab in Asia draws down its 45-day helium buffer at the same time. There is no routing around these constraints, no alternative supplier to call, no inventory to release. A single disruption at any one chokepoint propagates across the entire market within days, with no buffer to absorb it.

> **"The pace of industry supply growth could not keep up with demand, leading to broad-based price increases. Tight supply and demand conditions are expected to persist due to physical limitations in production space and ongoing technology migration."**
>
> **Song Hyun Jong, President & Head of Corporate Center, SK Hynix**
> Q4 2025 Earnings Call, January 29, 2026

AI Is Displacing Supply That Other Industries Depend On

AI infrastructure is not just growing demand in its own categories. It is pulling supply from markets that have no alternative source. The fab lines producing AI power components previously made automotive MOSFETs. The T-glass going into AI substrates previously went into smartphone packages. The helium prioritized for AI fabs is helium other manufacturing sectors cannot access. The transformers going to data centers are transformers that utilities and renewable energy projects cannot get. Those markets have no immediate path to recovery.

Category by Category

What's Short, By How Much, and What It's Costing

The data below is based entirely on supplier announcements, earnings disclosures, and confirmed distributor market data as of Q1 2026. None are modeled.

Component	Lead Time (Weeks)	Weeks
Power Transformers (Large Units)	Critical	128
Automotive eMMC	Critical	72
Ribbon Fiber (Data Center Grade)	Critical	64
High-Radix Network Switches	Critical	52
Power MOSFETs / Switching Regulators	Critical	52
800G Optical Transceivers	Elevated	48
Network ASICs (Tomahawk 5)	Elevated	44
Liquid Cooling CDUs	Elevated	36
T-Glass Substrates (High-Reliability)	Elevated	30
AI-Grade MLCCs	Moderate	20

Sources: Power Transformers — Wood Mackenzie, "Untangling the US Transformer Supply Chain Crisis," Q2 2025 • 800G Optical Transceivers, Network ASICs, High-Radix Switches, Liquid Cooling CDUs — Fusion Worldwide distributor market intelligence, Q1 2026 • Power MOSFETs / Switching Regulators — Infineon, Texas Instruments, ADI earnings disclosures, 2025–2026 • T-Glass Substrates — Goldman Sachs Research, T-Glass Shortage Analysis, 2025 • AI-Grade MLCCs — Murata, TDK, Samsung Electro-Mechanics channel data, Q1 2026 • Ribbon Fiber — Fusion Worldwide market intelligence, Q1 2026 • Automotive eMMC — Fusion Worldwide distributor market intelligence, Q1 2026

Memory & Storage

What Is Affected: HBM, DDR5 server DRAM, enterprise NAND, and SSD supply supporting AI data center deployments.

Lead Time & Pricing: HBM is fully allocated through 2026, with contracts locking in high-teens to low-20% price increases. As wafer capacity shifts toward HBM, traditional DRAM tightens behind it. That pressure is beginning to influence NAND availability, contributing to tighter SSD supply and rising pricing as hyperscalers consume larger volumes of enterprise storage.

Procurement Implication: This is no longer a single-product memory constraint. HBM prioritization is tightening DRAM, which is beginning to impact NAND and SSD availability. Procurement strategies focused only on compute risk underestimating the memory and storage layers required to make deployments operational.

"Supply cannot keep pace with demand, resulting in a severe supply and demand imbalance. Customer inventory levels have decreased overall, with server customers struggling to secure memory volumes."

SK Hynix
Q4 2025 Earnings Call, January 29, 2026

CPUs

What Is Affected: High-performance server CPUs supporting AI orchestration, data processing, inference, and emerging agentic AI workloads, including platforms from Intel and AMD.

Lead Time & Pricing: Lead time visibility varies by platform, but both Intel and AMD CPUs are experiencing extended lead times and delivery delays. Supply is increasingly constrained by advanced node availability, with AMD processors exposed to TSMC node tightness already under pressure from AI-related demand.

Procurement Implication: AI demand is expanding beyond accelerators. CPU requirements are increasing alongside system complexity, and agentic AI may drive higher-than-expected demand. If treated as a secondary consideration, CPU supply can become a secondary bottleneck in deployment timelines.

Advanced Packaging & Substrates

What Is Affected: CoWoS packaging slots, ABF substrates, T-glass reinforcement materials, copper-clad laminates (CCL), and OSAT capacity for leading-edge packaging operations.

Lead Time & Pricing: CoWoS: sold out through 2026. T-glass substrates (high-reliability grades): 28 to 32 weeks, cost up 22%. ABF substrates: 10 to 20% supply gap projected for 2026. CCL lead times: 16 to 20 weeks, suppliers declining to commit to firm dates.

Procurement Implication: Packaging slot availability determines whether a product ships. Without an established OSAT relationship or tier-one customer standing, an organization is last in the allocation queue regardless of price. At this stage, early commitment with volume guarantees is the only effective lever.

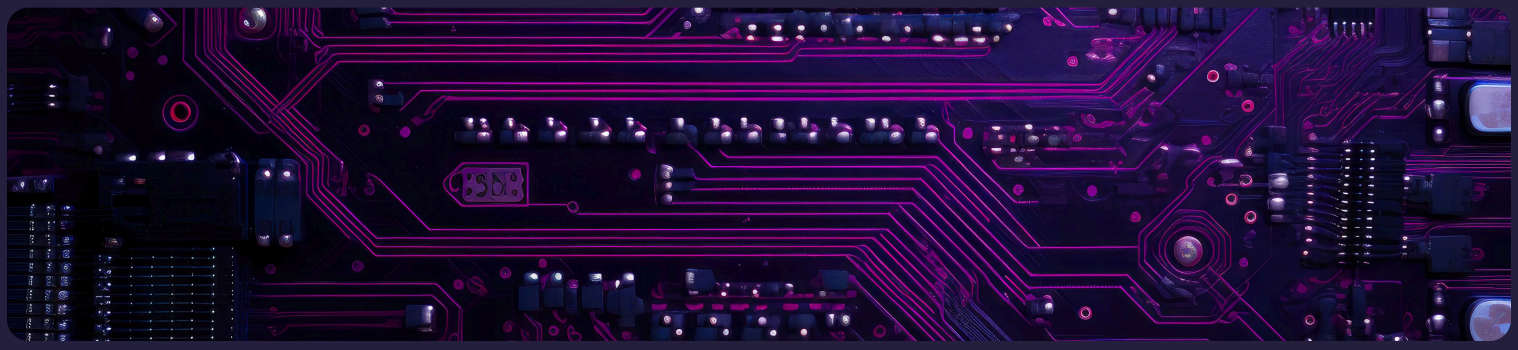

Material	2026 Constraint Summary
CoWoS (TSMC)	Sold out through mid-to-late 2026; NVIDIA holds ~70% of 2025 to 2026 capacity.
T-Glass (Nittobo)	20% price increase Aug 2025; new capacity not until late 2026; Apple competing for same supply.
ABF (Ajinomoto)	10 to 20% supply gap projected; customers signing multi-year LTAs to secure allocation.
CCL (T-Glass Grade)	16 to 20 week lead times; suppliers not committing to firm delivery dates.
ABF Substrate Film	Single commercial-scale producer. Tightening as AI packaging absorbs increasing share of output.

Power Management ICs & Passive Components

What Is Affected: Voltage regulator modules (VRMs), power MOSFETs, gate drivers, switching regulators, MLCCs, and tantalum capacitors.

Lead Time & Pricing: Low-voltage MOSFETs and switching regulators: 40 to 60 weeks. AI-grade MLCCs: less than one month of channel inventory remaining. Infineon, TI, and ADI have implemented 10 to 30% price increases. MLCC spot prices up 15 to 20% in Q1 2026.

Procurement Implication: Power management is the highest-volume, most consistently underestimated category on an AI BOM. AI server platforms require an order-of-magnitude more MLCCs per rack than traditional servers. At 40 to 60-week lead times, a mid-cycle re-order is a program delay. If you have not secured power component inventory against your full planned deployment volume, you are already behind the organizations that have.

~10x MLCC
Count Increase per AI Rack vs. Traditional Server

<1 Month
AI-Grade MLCC Channel Inventory Remaining, Q1 2026

40-60 Week
Lead Times
Power MOSFETs and Switching Regulators

128 Weeks
Average Lead Time Large Power Transformers
(Wood Mackenzie)

High-Speed Networking & Optical Components

What Is Affected: 800G and 1.6T optical transceivers, high-radix network switches, co-packaged optics, ribbon fiber, and networking ASICs.

Lead Time & Pricing: 800G transceivers: 36 to 50 weeks, pricing described as severe. Network switches: 40 to 56 weeks, up from a 12 to 16 week baseline. Ribbon fiber: 60+ weeks. Networking ASICs: 40 to 44 weeks.

Procurement Implication: Networking is not infrastructure that follows compute. Across sourcing activity we are tracking, networking components are still entering procurement cycles after compute allocation — creating delays even when accelerators and memory are secured. It is a parallel procurement dependency that must be committed on the same timeline as accelerators. A GB300 NVL576 mega-cluster requires 640+ switches versus roughly 12 for a 2022 HPC cluster. If networking is not in your procurement plan at the same time as compute, it will become the item that delays everything else.

Liquid Cooling & Thermal Management

What Is Affected: Coolant Distribution Units (CDUs), direct-to-chip cooling manifolds, pump assemblies, cold plates, and facility heat rejection infrastructure.

Lead Time & Pricing: CDU production is at capacity across every major vendor, with lead times extending well into 2026. The liquid cooling market is projected to grow from $2.8B today to over $21B by 2032. That growth rate means sustained supply pressure through the planning horizon, not a near-term correction.

Procurement Implication: Every GB300 NVL72 rack draws 132 to 142kW. It requires liquid cooling to operate. This is a physical prerequisite, not an optional upgrade. Organizations that treat thermal infrastructure as a late-stage procurement decision will find it is the item that holds up every other stage.

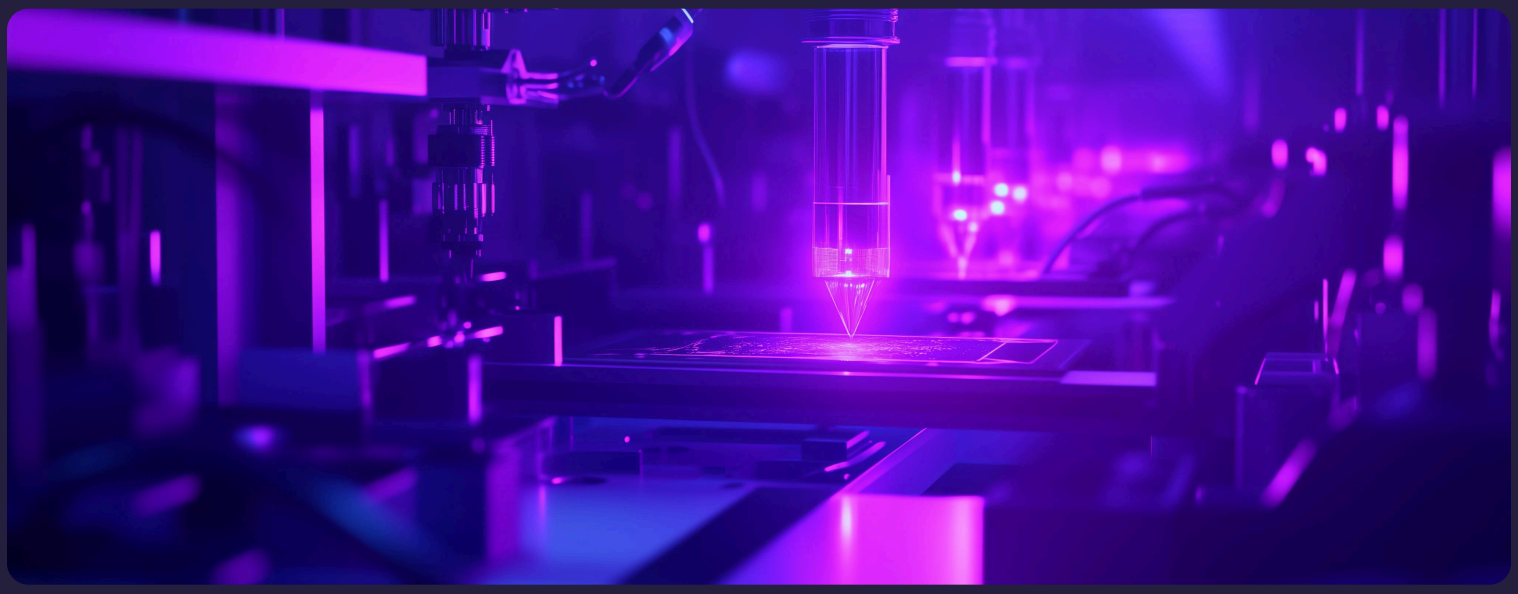

Power Transformers & Electrical Infrastructure

What Is Affected: Large power transformers, generator step-up units (GSUs), distribution transformers, switchgear, and high-voltage electrical infrastructure for data center and fab facilities.

Lead Time & Pricing: Power transformers: 128 weeks average, specialty fab units up to 200 weeks (Wood Mackenzie, Q2 2025). Prices up 77% since 2019. Major suppliers implementing 20 to 30% additional increases in 2026. Some manufacturers are declining to quote on large-scale projects.

Procurement Implication: At 128-week average lead times, a facility that has not already ordered primary power infrastructure for a 2027 go-live is behind. This category sits further upstream than most procurement teams track. It is also the one that, when it slips, holds up everything that follows.

Component Category	Price Movement
T-Glass Nittobo, Aug 2025	+20% blanket increase; ~$100/kg; orders backlogged into Q2 2026.
Power ICs Infineon	AI data center power revenue forecast +67% YoY; increases effective April 2026.
TI / ADI Power SKUs	+10 to 30% across thousands of legacy and power SKUs.
AI-Grade MLCCs	+15 to 20% spot price increase, Q1 2026; less than 1 month channel inventory.
800G Optical Transceivers	Described as severe by procurement managers; no published ceiling.
ABF / BT Substrates	Up to 20% increases anticipated; long-term contract premiums rising.
HBM3E Samsung, SK Hynix	+High-teens to low-20% on 2026 contracts; all 2026 production allocated.
Power Transformers Large Units	+77% since 2019; +20 to 30% from major suppliers in 2026.
Helium Spot Prices, March 2026	+70 to 100% since Ras Laffan outage; fabs drawing on 45-day buffer inventory.

What Leading Procurement Teams Are Doing to Stay Ahead in This Market

The approach that worked in 2021 to 2022, monitoring component pricing and waiting for lead time compression, does not work when the shortage is in consumed process inputs. There is no inventory correction coming. Correction requires new capacity, and new capacity takes years. Any procurement plan built on the assumption that this cycle corrects the way the last one did is a plan built on a false premise. The organizations outperforming are the ones that recognized this early and repositioned accordingly.

> **"The organizations deploying on schedule in 2026 moved earlier and saw further. Visibility into upstream material constraints is what made the difference. Speed alone cannot close a 12-month allocation gap."**

Andrew Czuczwa
Market Research Manager
Fusion Worldwide

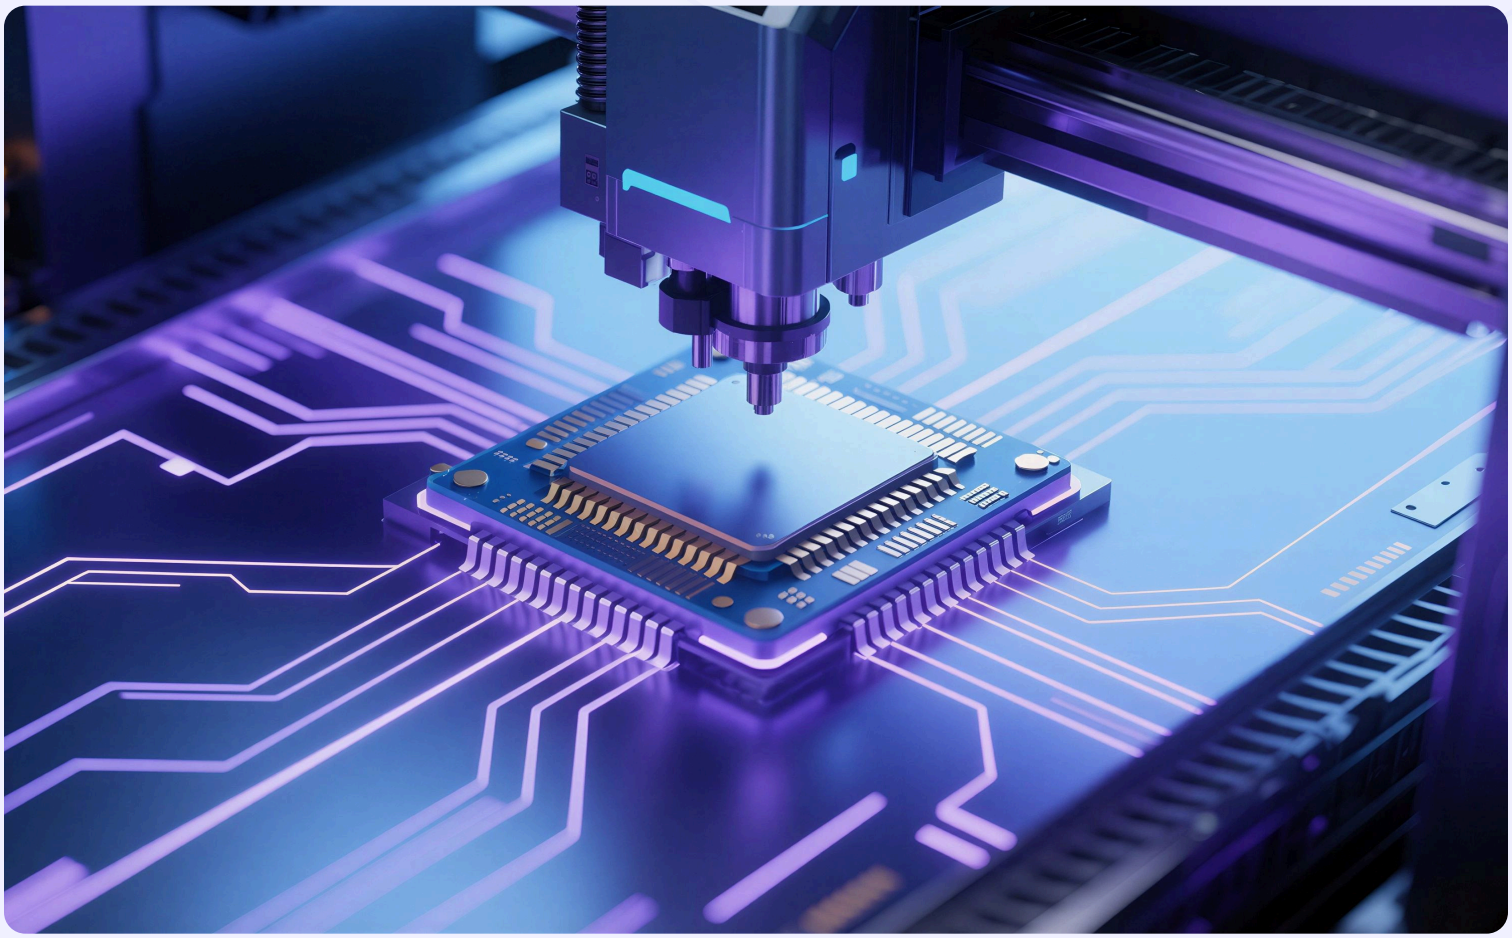

They know what's actually on their critical path.

At 40 to 128-week lead times, just-in-time procurement does not function. There is committed and uncommitted. Map the full AI infrastructure BOM against current lead times and identify every item above 24 weeks. That is the actual critical path. Any item above that threshold needs active management now, not at the next planning cycle.

They treat supplier relationships as supply chain assets.

In every constrained category in this report, allocation goes to the customers suppliers know, with commitments they can plan around. TSMC prioritizes long-term contracted buyers. T-glass supply goes to substrate manufacturers with existing agreements. CDU production slots go to buyers already in queue. HBM pricing and availability in 2027 is being settled in negotiations happening right now. The organizations that are in those negotiations have options. The ones that are not are buying at spot from whoever has remaining allocation, at premium pricing, from the end of the queue.

They're running market intelligence as a standing function.

The organizations that positioned ahead of this shortage were running active supply chain intelligence programs: tracking N3 wafer allocation windows, T-glass availability, helium spot prices, MLCC channel inventory, and OSAT booking windows in real time. From what we see across the market, the gap between organizations running active intelligence programs and those that are not is now measurable in deployment quarters, not weeks. Better information means better decisions, made earlier, at lower cost. It shows up directly in which organizations are deploying on schedule and which are not.

They've already built alternative source maps.

Single-source dependency is a risk that compounds as demand grows and primary allocation tightens. The organizations ahead of this curve have already mapped alternatives: OSAT providers including ASE Technology and Amkor for CoWoS-comparable packaging; Intel EMIB for hyperscaler ASIC packaging; Taiwan Glass under qualification as a second T-glass source; NVIDIA co-developing CoWoP with ASE SPIL. Having an alternative source qualified before you need it is a fundamentally different position than discovering you need one after primary allocation closes.

In 2026, the Bottleneck Is Everything Behind the Chip

The binding limit on AI infrastructure deployment is no longer the chip itself. It's the wafers the chip is etched on, the gases used to produce them, the fiberglass reinforcing its substrate, the films laminating its package, the power components managing its load, the optics connecting its cluster, the cooling systems keeping it alive, and the transformers delivering power to the building.

In this supercycle, constraints are distributed across the entire stack — and many sit upstream of where procurement teams are traditionally watching.

Patience is not a procurement strategy in this environment. The inputs now in shortage are consumed in production, not stockpiled for later release. Capacity expansion takes years, and the decisions shaping 2027 and the cycles that follow are being made now. In this cycle, constraints do not appear at the end of a build — they are embedded at the beginning. Organizations that recognize that early will deploy on schedule. The rest will discover the shortage only after it has already defined their timeline.

Leading procurement teams are widening sourcing strategies and engaging alternative supply before constraints fully surface. In a cycle where everything goes short, visibility and access matter more than ever. Fusion Worldwide helps organizations get **Out in Front™** of tightening supply — and stay there.

State of the Industry

FUSION WORLDWIDE

When Everything Goes Short, We Help You Keep Building

Search and buy components online, request quotes instantly, or work directly with a Fusion sourcing specialist. We connect procurement teams to global supply, alternative sources, and market intelligence to help you get **Out in Front™** — and stay there.

Contact Fusion Worldwide or visit fusionww.com to learn more.

Americas

Global Headquarters
Boston
+1 617 502 4100
boston@fusionww.com

EMEA

Regional Headquarters
Amsterdam
+31 20 667 6020
amsterdamoffice@fusionww.com

APAC

Regional Headquarters
Singapore
+65 6311 5250
singaporeoffice@fusionww.com

See which lead times are moving now.

Get Report