THE PHYSICAL WORLD

UNDERNEATH AI

Five Supply Chain Pressure Points That Will Define the AI Buildout in 2026

February 2026

The Trillion-Dollar Traffic Jam

There is a story the market tells about artificial intelligence. It goes like this: the smartest companies design the most powerful chips, train the largest models, and whoever has the best software wins. It is a clean story. It fits on a slide deck. And it is dangerously incomplete.

The real story of AI in 2026 is a story about physics. About heat, about copper, about the speed of light through glass, and about the fact that a single GPU package now consumes more power than an American household. It is the story of what happens when an industry that has spent decades abstracting away the physical world suddenly slams into the limits of that physical world at full speed.

This document traces the five critical bottlenecks in the AI infrastructure supply chain as they stand in early 2026. They are not abstract risks. They are concrete, measurable constraints — the places where the entire AI buildout either accelerates or stalls. Understanding them is not optional for anyone with capital at risk in this cycle.

The five are: advanced semiconductor packaging, high-bandwidth memory, the transition from copper to optical interconnects, the liquid cooling supply chain, and the coming shift to 800-volt DC power architecture. Each one is a story in itself. Together, they form the map of what it actually takes to build the physical world underneath AI.

Chapter One: The Packaging Chokepoint

Why the most advanced chips in the world are stuck in a queue

To understand the first and most fundamental bottleneck in the AI supply chain, you need to understand a counterintuitive fact: manufacturing a cutting-edge AI chip is no longer the hardest part. TSMC can fabricate billions of transistors on a 3-nanometer wafer with remarkable yield. The hard part is what comes after.

Modern AI processors — Nvidia's Blackwell, AMD's MI300, Google's TPUs — are not single chips. They are systems of chiplets: multiple silicon dies from different vendors, stacked and stitched together on an advanced substrate using a technology called Chip-on-Wafer-on-Substrate, or CoWoS. The GPU die, the high-bandwidth memory stacks, the I/O chips — they all need to be co-packaged with micron-level precision. Without this packaging step, even a perfectly fabricated 3nm wafer is just an expensive coaster.

And this is where the entire industry hits a wall. TSMC's CoWoS capacity has been sold out through 2025 and deep into 2026. The company is ramping from roughly 75,000 CoWoS wafers per month toward a projected 120,000 by the end of 2026. But demand is doubling annually, which means the gap between what the market wants and what can physically be built is persistent and growing.

The numbers tell the story: chiplets — the heterogeneous dies from different vendors that get co-packaged together — are estimated to reach $100–110 billion in annual revenue in 2026. This is not a niche market. This is the market. And the capacity to package these chiplets into finished products sits in an alarmingly small number of facilities, predominantly in Taiwan.

The second-tier packaging houses, known as OSATs (Outsourced Semiconductor Assembly and Test), are scrambling to build capability. Amkor and ASE are investing heavily. Intel's foundry services division is positioning its own advanced packaging as an alternative. Samsung has expanded its HBM packaging lines. But the learning curves are steep, yields matter enormously when you are bonding $10,000 worth of silicon onto a single substrate, and TSMC's head start is measured in years, not months.

Here is the key insight for anyone thinking about the AI supply chain: CoWoS is the binding constraint. It sits upstream of everything else. You cannot sell an AI GPU that hasn't been packaged. You cannot deploy a server that doesn't have GPUs. Every other bottleneck we will discuss — memory, optics, cooling, power — only matters if the chips themselves can be

assembled and shipped. Advanced packaging is the toll booth through which the entire AI buildout must pass.

Chapter Two: The Memory Supercycle

Why this time really is different for DRAM

Memory has always been a cyclical business. The semiconductor industry has lived through decades of boom-and-bust cycles — oversupply crashes followed by undersupply spikes — and investors have been conditioned to treat any memory rally with suspicion. The current HBM (High Bandwidth Memory) supercycle is asking them to unlearn that instinct.

Here is what makes HBM structurally different from ordinary DRAM. A standard DDR5 memory module is a commodity. It goes into laptops, servers, phones. Demand is broad, supply can be ramped relatively quickly, and pricing is set by competitive dynamics among three oligopolists — Samsung, SK Hynix, and Micron. HBM is none of those things.

HBM is a vertically stacked tower of DRAM dies, bonded together with through-silicon vias (TSVs), then integrated onto the same advanced package as the GPU. Current-generation HBM3E stacks eight or twelve dies high. The manufacturing process is extraordinarily complex: each die must be thinned to roughly 30 microns (thinner than a human hair), perfectly aligned, and bonded with near-zero defects. A single bad layer can ruin an entire stack.

The result is a product with bandwidth that no other memory technology can touch. HBM3E delivers roughly 8 terabytes per second of bandwidth per GPU package. The upcoming HBM4, targeted for Nvidia's Vera Rubin platform, is designed to hit 22 TB/s — nearly three times what Blackwell offers — with 288GB of capacity per socket.

Both Micron and SK Hynix have their HBM capacity sold out through calendar 2026. SK Hynix maintains more than 50% market share in HBM and is expected to hold that position through at least 2026. The total HBM market is projected to grow from roughly $35 billion in 2025 to around $100 billion by 2028. Revenue share of HBM is expected to surpass 30% of all DRAM revenue by 2026, despite HBM representing only about 8% of total DRAM wafer output.

This creates a fascinating secondary effect that most analysts overlook: every DRAM wafer dedicated to HBM is a wafer not producing conventional memory. The pivot to HBM is cannibalizing supply of the memory that goes into PCs, smartphones, and traditional servers. If HBM demand continues to surge — and every indication says it will — conventional DRAM prices will rise too, not because of demand for those products, but because of supply displacement. The three-company oligopoly has both the incentive and the ability to manage this dynamic to their advantage.

The transition from HBM3E to HBM4 is happening simultaneously with all of this. HBM4 moves from a 1024-bit interface to a 2048-bit interface, introduces a logic base die manufactured on advanced process nodes, and requires even tighter integration with the GPU package. It is not simply “more of the same.” It is a generational leap in manufacturing complexity, arriving at precisely the moment when existing capacity is already maxed out.

Chapter Three: The Speed of Light

How co-packaged optics will rewire AI datacenters

If advanced packaging is the first bottleneck and memory is the second, the third is one that fewer people have on their radar — but it may be the most structurally important over the next three to five years. It is the transition from copper to optical interconnects inside the datacenter.

Here is the problem in plain terms. An AI training cluster today might connect tens of thousands of GPUs using high-speed cables. Those cables have traditionally been copper. Copper works fine at moderate speeds over short distances. But the bandwidth demands of modern AI clusters have pushed copper to its physical limits.

At terabit-per-second speeds, copper cables become thicker, stiffer, and shorter-range. The signal degrades over distance. You need active components to boost the signal. The cables themselves become so unwieldy that they create physical density problems inside the rack — you literally cannot fit enough copper cables into the space to connect everything that needs to be connected. Power consumption per link is high, and it scales poorly.

The pluggable transceiver era

The industry's first answer to this problem was pluggable optical transceivers — small modules that convert electrical signals to light and back. These plug into the front panel of a switch and connect via fiber optic cables. They work, and they have been the backbone of datacenter networking for years. Companies like Coherent, Lumentum, and Innolight have built billion-dollar businesses around them.

But pluggable transceivers have their own limitations. At 800 gigabit and 1.6 terabit speeds, the electrical signal must travel from the switch ASIC, across the circuit board, through a connector, and into the transceiver module before being converted to light. That journey is lossy — about 22 decibels of signal loss, which must be compensated for with power-hungry electronics. A single pluggable transceiver port at 1.6T consumes around 30 watts.

Co-packaged optics: putting the laser next to the chip

Co-packaged optics, or CPO, takes a fundamentally different approach. Instead of plugging optical modules into the front panel of a switch, CPO integrates optical engines directly onto the switch package itself — right next to the ASIC, on the same substrate. The electrical signal only needs to travel millimeters before being converted to light, rather than centimeters or inches.

The performance improvement is dramatic. Signal loss drops from 22 dB to about 4 dB. Per-port power consumption falls from 30 watts to approximately 9 watts. To understand why this matters at scale, consider a large AI cluster with 2.4 million links. With pluggable transceivers, those links consume roughly 72 megawatts of power just for the interconnect. With CPO, that drops to about 21.6 megawatts — a 3.3x reduction. At electricity prices of $0.05–0.10 per kilowatt-hour, that is tens of millions of dollars per year in operating savings, before you even count the cooling savings from dissipating less heat.

Cost is compelling too. CPO optical engines are estimated at $35–40K per switch equivalent, compared to roughly $72K for equivalent pluggable transceivers. And because CPO eliminates the front-panel connector, it frees up physical space that can be used for additional ports or denser configurations.

2026: the year CPO enters production

This is not theoretical. Broadcom's Tomahawk 5-Bailly is already shipping as the industry's first volume-production CPO switch. Nvidia's Quantum-X InfiniBand Photonics is arriving in early 2026, with Spectrum-X Ethernet Photonics following in the second half. Both are built on TSMC's COUPE packaging platform, which itself is a derivative of the same CoWoS advanced packaging that gates GPU production.

A critical design choice in CPO deserves attention: the external laser source. Lasers are the most failure-prone component in any optical system. CPO architectures have standardized on keeping the laser separate from the optical engine, housed in a field-replaceable module. When a laser degrades — and they do, over thousands of hours of operation — you can swap the laser without replacing the entire switch package. This is an elegant solution to a reliability problem that could otherwise make CPO adoption risky.

The market is projected to exceed $20 billion by 2036, growing at a 37% compound annual rate from 2026. But the more important number is this: the scale-up fabric inside AI clusters — the connections between GPUs within a training run — is currently 100% electrical. Converting even a fraction of that to optical represents an opportunity that is orders of magnitude larger than the scale-out networking market where pluggable transceivers live today.

The CPO supply chain map

What makes CPO particularly interesting from an investment perspective is that the supply chain is still being built. Unlike GPUs or memory, where the competitive dynamics are well-understood, CPO touches a different set of companies across multiple geographies.

Taiwan has the most complete CPO supply chain outside the United States. TSMC provides the COUPE packaging platform — the foundation on which both Nvidia and Broadcom build their CPO switches. ASE, the world's largest outsourced semiconductor assembly company, provides the advanced assembly capability. Accton Technology builds white-box switches and is mass-producing 800G to 1.6T silicon photonics modules. Quanta Computer, as a major ODM, builds the rack-scale systems. Taiwan has even formed a "Silicon Photonics Industry Alliance" (SiPhIA) with TSMC, ASE, MediaTek, and Quanta as members, targeting what they estimate could be a $100+ billion AI inference market opportunity in the 2026–2028 boom period.

In the United States, the switch ASIC layer is dominated by Nvidia and Broadcom, with Marvell making a bold entry through its $5.5 billion acquisition of Celestial AI in December 2025 — a bet that optical interconnects become standard plumbing for custom AI accelerators. On the laser and component side, Coherent is the vertically integrated leader, supplying continuous-wave lasers, external laser sources, and detectors. The company has secured what it described as an "exceptionally large" CPO purchase order from a major AI datacenter customer, with initial revenue expected at the end of calendar 2026. Coherent's most recent quarter showed $1.7 billion in revenue, with the datacenter and communications segment at $1.2 billion — up 33.6% year-over-year. Lumentum is the other major laser house, and Nvidia has strategically locked up its EML (electro-absorption modulated laser) capacity, pushing lead times past 2027.

In Japan, Sumitomo Electric Industries offers end-to-end capability from optical fiber to connectors to transceivers, while Furukawa Electric has a similar profile. In Korea, the exposure is more indirect — through HBM packaging and substrates that sit alongside CPO optical engines in the same CoWoS packages.

For GlobalFoundries, the acquisition of Singapore's Advanced Micro Foundry in November 2025 gave it arguably the most important dedicated silicon photonics foundry outside TSMC. Tower Semiconductor in Israel is investing $650 million to triple its silicon photonics shipments by mid-2026.

Crucially, CPO does not kill the pluggable transceiver business — it opens a new total addressable market. The scale-out networking that pluggable transceivers serve continues to grow. CPO's opportunity is in converting the scale-up fabric from electrical to optical, a market that barely exists today but could dwarf the existing transceiver market by the end of the decade.

Chapter Four: The Heat Problem

Liquid cooling and the barbell bet on Vertiv and Modine

In 2023, a standard datacenter rack consumed 10 to 15 kilowatts of power. By early 2026, AI-specific clusters are pushing 120 to 150 kilowatts per rack. By 2027, Nvidia's Kyber rack architecture will pack 576 Rubin Ultra GPUs into a single rack consuming 600 kilowatts — a 40x increase in three years.

Air cooling stopped being viable somewhere around 40–50 kilowatts per rack. Beyond that threshold, no amount of fan speed or airflow engineering can remove heat fast enough to prevent throttling. The physics simply do not work. This is why direct-to-chip liquid cooling has become the default for every serious AI deployment, commanding roughly 65% of the liquid cooling market in 2026.

The liquid cooling market is growing from $5.1 billion in 2025 to a projected $16 billion by 2030. But the interesting story is not the market size — it is the component-level bottlenecks within the cooling supply chain that determine who wins.

The cooling stack

Think of liquid cooling not as a single product but as a stack of interdependent components, each with its own competitive dynamics.

Coolant Distribution Units (CDUs) are the bottleneck. These are the industrial heat exchangers that sit between the IT cooling loop and the facility water system. Lead times are 12 to 18 months. CDU revenue is expanding roughly 6x, from about $150 million in 2023 to a projected $1 billion by 2027. Even customers who have signed hyperscaler contracts and secured power allocations cannot deploy AI infrastructure without CDU equipment. This is the constraint that keeps datacenter operators up at night.

Cold plates are the component that sits directly on the GPU die, transferring heat into the liquid coolant. Manufacturing requires friction stir welding or vacuum brazing of copper micro-channels, followed by 100% leak testing. Revenue is growing from under $200 million to roughly $1.4 billion between 2023 and 2027. This is becoming increasingly competitive as more manufacturers build capability, meaning margins may compress over time.

Quick-disconnect fittings are the hidden monopoly in the stack. Danfoss, a Danish company, supplies the majority of quick-disconnect connectors for AI datacenters and is Nvidia's recommended vendor. These are small, inexpensive components — but the

switching cost is enormous because any leak in a liquid cooling system can destroy tens of millions of dollars worth of GPU hardware. Customers are deeply reluctant to change connector vendors. Parker-Hannifin and Stäubli are alternatives, but inertia favors the incumbent.

Fluid management and services represent the operational layer. Liquid cooling systems are not plug-and-forget. They face real failure modes: corrosion from mixed metals in the loop, contamination during installation from incompatible materials, and microbial growth in idle coolant lines. These are unglamorous problems, but they determine whether a $200 million AI cluster actually works or slowly degrades.

The barbell: Vertiv and Modine

Vertiv Holdings has emerged as the consensus blue-chip play on datacenter cooling. The numbers are impressive: 52% stock gain in 2025, $9.5 billion backlog, 28% organic sales growth, over 80% of revenue tied to datacenters. Vertiv co-develops reference architectures with Nvidia for each GPU generation, gives it a first-mover advantage on specifications. Its 2025 acquisition of PurgeRite for $1 billion — a specialist in mechanical flushing and fluid management — was strategically smart, giving the company ownership of the entire "liquid lifecycle" from deployment through ongoing maintenance.

But at roughly $71 billion in market capitalization and approximately 68x trailing P/E, Vertiv trades at a premium that already reflects dominant market position. Schneider Electric and Vertiv are in a dead heat for global datacenter infrastructure market share as of 2026. And hyperscale customers are actively diversifying their cooling vendor base — nobody wants single-supplier risk on something this critical to operations.

This is where Modine Manufacturing becomes interesting as the asymmetric counterpart. Modine offers near-pure-play AI cooling exposure at roughly half Vertiv's P/E multiple. The company, founded in 1916 in Racine, Wisconsin, transformed itself through its acquisition of Airedale, a UK-based precision cooling specialist. In early 2025, Modine announced a $180 million order — the largest in the company's 109-year history — from a hyperscaler. Management is guiding for 60%+ datacenter revenue growth for the full fiscal year 2026, targeting over $2 billion in datacenter revenue by FY2028.

The honest framework is a barbell: Vertiv for the execution story at the blue-chip end, Modine for the asymmetric upside at a lower multiple. Both are "capex derivative" businesses — if hyperscaler spending slows for any reason, the entire cooling stack gets repriced simultaneously. But as long as the buildout continues, the physics-driven demand is as close to non-discretionary as infrastructure spending gets.

Two additional names deserve mention. nVent Electric has roughly 30% datacenter exposure and a growing Nvidia partnership that is underappreciated by the market. And Eaton’s $9.5 billion acquisition of Boyd for liquid cooling gives it instant scale, though at a premium price.

Chapter Five: The Voltage Revolution

Why 800-volt DC power is the next infrastructure transition

The final piece of the infrastructure puzzle is the one that ties everything else together. It starts with a number that sounds mundane until you do the math: in a rack consuming one megawatt of power at the industry-standard 54 volts DC, the copper busbar required to carry that current weighs up to 200 kilograms. In a one-gigawatt datacenter with a thousand such racks, the busbars alone require 200,000 kilograms of copper. That is 200 metric tons of copper just for internal power distribution — before you count the wiring, the transformers, or any other electrical infrastructure.

This is not economically sustainable. It is not even physically practical. And it is why Nvidia is leading the industry toward 800-volt DC power distribution for its next-generation platforms.

The physics are straightforward: power equals voltage times current. For a given power delivery requirement, doubling the voltage halves the current, which dramatically reduces the conductor cross-section needed. Moving from 48V or 54V to 800V DC reduces copper requirements by roughly 45%, improves power efficiency by up to 5x compared to conventional 48V methods, and eliminates multiple AC-to-DC conversion stages that each waste energy as heat.

The electric vehicle and solar industries have already made this transition — 800V DC architectures are standard in modern EVs. Datacenters are following the same path, driven by the same physics.

The Nvidia roadmap: Oberon to Kyber

To understand the 800V timeline, you need to understand Nvidia's platform roadmap.

The Vera Rubin NVL72 and NVL144, arriving in the second half of 2026, are designed to fit within the existing Oberon rack architecture at 120–130 kilowatts per rack. This is backward-compatible with today's datacenter infrastructure. It is the "easy" deployment.

Everything changes with Rubin Ultra and the Kyber rack in 2027. The Kyber rack rotates compute blades 90 degrees into a vertical orientation, packing 576 Rubin Ultra GPUs into a single rack. Each Rubin Ultra GPU package consumes approximately 3.6 kilowatts and includes up to 1 terabyte of HBM4E memory at 32 TB/s bandwidth. The total rack power draw hits 600 kilowatts — five times current systems.

Six hundred kilowatts cannot be delivered at 48 volts. The current required would melt conventional busbars. This is why the Kyber rack is designed from the ground up for 800V DC, and why Nvidia announced the specification alongside more than 20 industry partners at OCP Summit.

The 800V supply chain

Three components define the 800V DC supply chain.

Solid-state transformers (SSTs) replace traditional line-frequency transformers, converting medium-voltage AC directly to 800V DC. This market is projected to grow at a 32% compound annual rate, reaching roughly $1 billion by 2030. The key players are the industrial power giants: Eaton, GE Vernova, Schneider Electric, Siemens, and Delta Electronics. These are companies with decades of power conversion experience, and their engineering teams are designing the specific SST topologies needed for datacenter deployment right now.

GaN (gallium nitride) semiconductors are critical for the high-efficiency DC-DC converters that step 800V down to the low voltages that GPU cores actually require. GaN transistors switch faster and more efficiently than traditional silicon MOSFETs, which matters enormously when you are converting hundreds of kilowatts per rack. But the GaN supply chain is tight — notably, TSMC closed its GaN production line, creating material scarcity. Key vendors include Navitas Semiconductor, Infineon, STMicroelectronics, and Innoscience.

Safety and workforce requirements are the underappreciated third component. 800V DC is not something a standard datacenter technician can work on. It requires “qualified persons” with specific electrical safety training, new PPE protocols, arc flash analysis, and redesigned maintenance procedures. This is a significant operational change for an industry accustomed to low-voltage rack work.

Vertiv has announced a complete 800V DC product line for 2026 supporting the Nvidia Rubin Ultra platform, positioning itself at the intersection of cooling and power delivery. Foxconn is building its Kaohsiung-1 datacenter in Taiwan specifically designed for 800V DC from the ground up. CoreWeave, Lambda, Nebius, Oracle Cloud Infrastructure, and Together AI are among the early adopters designing 800V-native facilities.

The timeline is compressed: these components need to be designed, validated, and in supply chains now for 2027 deployment. GaN semiconductors, solid-state transformers, new

busbar designs, and 800V-rated connectors and protection equipment all need to be secured in 2026.

The Map of the Physical World

Step back and look at these five constraints together, and a picture emerges.

At the top of the stack sits advanced packaging — the binding constraint that gates everything downstream. If CoWoS capacity does not expand fast enough, nothing else matters. Below that is memory, where the HBM supercycle is structural rather than cyclical, driven by physics rather than sentiment, and controlled by an oligopoly with pricing power.

In the middle of the stack, the copper-to-optics transition via CPO is the most architecturally significant change, the one with the longest runway, and the one where the supply chain map is still being drawn. It is early enough that the competitive dynamics are unclear, which is usually where the edge is for investors willing to do the work.

At the physical layer, liquid cooling has gone from optional to non-negotiable, creating a new industrial supply chain where CDU capacity is the binding constraint and a barbell strategy — Vertiv at the blue-chip end, Modine at the asymmetric end — captures the opportunity across the risk spectrum.

And underpinning all of it, the 800V DC power transition is the change that makes everything else possible at gigawatt scale. Without it, the racks that Nvidia is designing for 2027 simply cannot be built. It is the least discussed of the five constraints and arguably the one with the most certain demand trajectory.

The common thread is that every one of these bottlenecks is a physics problem, not a software problem. You cannot prompt-engineer your way to more CoWoS capacity. You cannot train a model to make copper carry more bandwidth. You cannot fine-tune gravity to help heat rise faster.

The companies that solve these physical problems — the ones that can actually manufacture, package, connect, cool, and power the hardware that AI runs on — are the infrastructure layer of the most important technology transition of the decade. The market is beginning to understand this. The question for investors is whether the understanding has arrived faster than the capital.

To view holdings of AIS, click here: https://www.vistashares.com/etf/ais/
