

The Architecture of Disruption

Horacio Coutiño
Multi-Asset Strategist
hc@exante.eu

"The test of a first-rate intelligence is the ability to hold two opposed ideas in mind at the same time and still retain the ability to function."

— F. Scott Fitzgerald

The AI buildout is no longer an abstract software story, but a capital-intensive story that is reshaping bond markets, power grids and equity leadership across the cycle. It begins with a physical footprint that has few historical precedents: hyperscalers have committed CapEx to data centres whose power density, design churn and cost per megawatt are rising fast enough to turn yesterday's 'long-lived' infrastructure into tomorrow's stranded assets. That buildout is being financed through an increasingly complex architecture of on-balance-sheet bond issuance and off-balance-sheet shadow borrowing, a structure that is quietly expanding investment-grade duration supply and steepening the long end of the US yield curve.

At the same time, the AI capex super-cycle is colliding with geopolitical and commodity constraints. The Iran shock has exposed how vulnerable the ecosystem is to LNG, helium and tungsten bottlenecks, turning energy prices and critical minerals into binding constraints on semiconductor output just as GPU demand inflects.

Against this backdrop, this analysis organises the AI taxonomy into seven investable layers: fabrication tools and process materials, memory, compute, networking, physical infrastructure, power and applications. It then builds concentrated baskets of listed companies that best capture each layer's signal. Using forward EPS, EBITDA, FCF margins and EV/EBITDA, it shows that earnings power and valuation have re-rated unevenly across the buildout.

The physical footprint of AI that defies casual analysis

Before examining its second-order effects, it is worth setting out the scale of the build-out clearly. The five largest US hyperscalers - Amazon, Alphabet, Meta, Microsoft and Oracle - have collectively committed between $660 billion and $690 billion in capital expenditure in 2026 alone, nearly doubling 2025 levels. BloombergNEF reports that the CapEx of the 14 largest publicly owned data centre operators globally approaches $750 billion this year. Goldman Sachs Global Institute's analysis, anchored to NVIDIA forward revenue estimates, projects cumulative capital requirements of approximately $7.6 trillion between 2026 and 2031 across compute, data centres and power infrastructure. The baseline model implies $765 billion in annual AI CapEx in 2026, growing to $1.6 trillion in annual CapEx in 2031.

The physical anatomy of what is being built is now more complex and more expensive than most market commentary acknowledges. The core unit is the AI accelerator and large-scale data centres house hundreds of thousands of these chips. Today's leading configurations, such as NVIDIA'S GB300 NVL72, pack 72 processors into a single rack consuming upward of 130 to 200 kilowatts of power. As AI workloads increase power density and require deeper system integration, the cost of building data centres has risen materially compared with prior generations of cloud infrastructure. These design changes have translated directly into higher capital costs per megawatt. Traditional hyperscale cloud facilities were typically built at approximately $10 million per megawatt (MW). However, current market commentary suggests that next-generation AI data centres are increasingly being developed at costs of $15 million to $20 million per MW. This range may rise further as power density and redundancy requirements continue to increase. The long-lived asset assumption that anchored prior data centre valuations can increasingly become a value trap when the design requirements of a facility shift materially within a few years of commissioning.

As of end-2025, according to BloombergNEF, over 23.1 gigawatts (GW) of data centre capacity was under construction globally, split across 831 sites, with 17 GW of this capacity being built in the Americas region, across 311 locations. The IEA projects global data centre electricity consumption will reach 945 terawatt-hours (TWh) in 2030, more than 2x the 415 TWh consumed in 2024. The US and China together will account for nearly 80 percent of global growth in that figure to 2030. None of these projections existed in their current form three years ago. The scale has compounded faster than the institutions built to regulate, finance or power it.

Funding the AI frontier and yield curve dislocation

The most underappreciated structural dynamics of the current cycle is not the scale of capital being deployed. It is the architecture through which that capital is being raised and the secondary consequences of that architecture for fixed income markets, term premia and the interest rate environment on which the buildout itself depends.

For most of the past decade, hyperscalers largely ignored the public bond market. Their cash generation was extraordinary, their balance sheets were underleveraged by any historical standard and debt financing was simply unnecessary. In 2025, the five largest US hyperscalers (Amazon, Alphabet, Meta, Microsoft, and Oracle) issued $121 billion in investment-grade corporate bonds, compared to an annual average of $28 billion between 2020 and 2024. By mid-March 2026, the group had already issued more debt than during the entirety of 2025. By October 2025, the total stock of AI-related investment-grade debt had increased to $1.2 trillion, making it the largest segment in the investment-grade market, surpassing US banks at 14% of the JP Morgan US Liquid Index. The Wall Street Journal reported that the trillions required to fund data centres could force tech hyperscalers to expand the corporate bond market by as much as twenty percent per year. S&P Global Ratings, in its 2026 Credit Liquidity Outlook, noted that tech sector debt issuance reached

record highs in 2025, representing 16.7% of global non-financial corporate bond issuance, up from 11.6% just twelve months prior.

However, the publicly visible bond issuance is just the tip of the iceberg. The Bank of International Settlements (BIS), in its March 2026 Quarterly Review, documented a parallel and rapidly expanding universe of off-balance-sheet financing that it termed 'shadow borrowing'. The architecture is consistent. A special purpose entity or joint venture, capitalised with equity from a consortium of sponsors, raises debt through private placements to develop data centre assets. The hyperscaler holds a minority stake, commits to long-term operating leases or take-or-pay capacity agreements and may provide guarantees against the vehicle's obligations. Economically, CapEx is converted to OpEx; the debt sits off the hyperscaler's balance sheet, but is guaranteed by it. The risk is transferred to private credit funds, insurance companies and institutional investors who hold instruments backed by contractual cash flows from investment-grade counterparties. The Meta and Blue Owl Capital $27 billion Special Purpose Vehicle (SPV) financing Meta's Louisiana AI data centre campus, and Oracle's $38 billion project-finance loan for Texas and Wisconsin AI facilities, are among the most visible expressions of this model.

BIS officials Egemen Eren, Ingomar Krohn and Karamfil Todorov made the systemic implications explicit: 'These structures strengthen links between hyperscalers and non-bank investors such as private credit vehicles and insurers. Banks support the vehicles with funding lines, potentially creating new shock transmission channels.'

The implications for bond markets are structural. The Dallas Fed identifies three distinct channels through which this issuance affects duration supply in US interest rate markets. The mechanisms matter because they operate differently on swap spreads versus Treasury yields, with implications that are still working their way through positioning.

The first mechanism operates through duration supply: most hyperscaler issuance has been longer-term, with maturities exceeding five years, locking in funding for multi-year build-outs. Longer-term investment-grade corporate bond issuance by hyperscalers adds duration supply to the market. The sheer volume of supply arriving in a compressed window is a mechanical pressure on long-end rates that operates independently of the Fed's posture.

The second channel is subtler and concerns the swapping of floating-rate loans from private credit investors into fixed-rate exposure. AI infrastructure financing has penetrated private credit markets in force. According to Bloomberg, convertible bond issuance linked to AI companies surged more than 556 percent y/o/y in the first weeks of 2026 and private credit vehicles have written deals against future recurring revenue streams and AI monetisation potential. When these floating-rate private credit loans are swapped into fixed-rate obligations, the resulting synthetic duration supply should bias Treasuries richer versus swaps. It pressures swap yields higher relative to Treasury yields, creating the kind of divergence between the long end of the yield curve and long-term swap spreads that the Dallas Fed identifies as the empirical signature of these flows.

The third channel is crowding out. The Big Five hyperscalers' $140 billion annual borrowing pace effectively rivals the Big Six banks, which have been the dominant issuers in the IG market. If hyperscaler supply crowds out financial issuers rather than simply adding to aggregate supply, the mere substitution of technology issuance for financials, who characteristically issue shorter duration, translates into incremental duration supply to the market as a whole, pressuring the long end even if overall issuance quantities remain unchanged.

In some respects, these companies are seen as more creditworthy than sovereign issuers, Alphabet, for example, sold a 100-year sterling bond with a 6.125% coupon that drew a bid-

to-cover ratio above 10.0x. Yet the sheer volume of tech issuance has created a market dislocation as the bonds of these high-quality issuers can be bought at a spread pick up over weaker fundamentals.

The more structurally concerning implication, identified by the Dallas Fed's Srini Ramaswamy, Seth Searls and Hugo De Vere, is that the volume and term structure of AI-related debt issuance is now a persistent upward pressure on long-end US interest rates. The Fed can manage the overnight rate. It cannot directly manage the duration supply being generated by a multi-year infrastructure super-cycle. The divergence between short and long rates means a steepening impulse is likely to persist and may intensify as the build-out accelerates.

In addition, higher-for-longer rates raise the weighted average cost of capital for AI infrastructure projects whose return horizons extend over years rather than quarters. They increase the debt service burden of the hyperscalers. Credit markets have remained remarkably receptive to issuance volumes so far, but receptiveness is not the same as equilibrium; it reflects demand that has not yet been tested against an inflection point in either credit quality or yield level.

The off-balance-sheet structures introduce a second-order systemic concern that the BIS has explicitly flagged. The 2008 analogy is imperfect, but not irrelevant: Structured Investment Vehicles expanded off-balance sheet and were abruptly brought onto balance sheets during the stress event, crystallising liabilities that had been invisible to markets. If the AI infrastructure vehicles face refinancing stress, whether from deteriorating private credit appetite, covenant violations or the activation of hyperscaler guarantees, the resulting liability migration onto corporate balance sheets would arrive as AI revenue realisation is under scrutiny. The question is whether these structures remain sound under scenarios of sustained higher rates, delayed monetisation or a material deterioration in the buildout narrative. AI revenues are still a fraction of CapEx commitments. The structures are being underwritten on the expectation of monopoly-like infrastructure rents from AI dominance. That expectation might not be irrational, but it is not yet a fact.

Geopolitical risk and the critical mineral squeeze on AI CapEx

The re-escalation of tensions with Iran in late February 2026 introduced the most acute stress test that cycle has faced yet. It operates through two mutually reinforcing channels — one financial, one industrial — neither of which the market has fully priced in.

The financial transmission is well understood in structure, if not yet in magnitude. An energy-driven inflation shock complicates the rate outlook at precisely the moment when AI CapEx is most exposed to it. Fed rate cut expectations faded entirely following the initial strikes. ECB deposit rate expectations swung from easing to tightening within a single trading day. Bonds sold off, not rallied, because investors priced the inflationary consequences of the oil shock ahead of safe-haven demand. This represents a continuation of the bond-equity correlation breakdown documented in our Alpha Vibes edition, The Strait of Reckoning: the 3-month bond-equity correlation for Germany reached 0.77, Japan 0.27, and the US 0.23. These were dramatic reversals from the deeply negative readings that prevailed through much of 2025. Allianz Research confirms the historical link: the correlation between energy price volatility and US Information technology sector CapEx is among the strongest across all S&P 500 sectors since 2000. Energy shocks reliably translate into investment hesitation across the tech supply chain. Higher-for-longer rates give investors and lenders legitimate cover to tighten standards, demand stronger return guarantees and quietly stall the most speculative infrastructure projects at the moment when the CapEx-monetisation gap is under scrutiny.

The industrial pressure point is more concrete and more immediate. Qatar's Ras Laffan complex, the world's largest LNG liquefaction facility and responsible for one-fifth of global LNG supply, sustained missile damage that QatarEnergy's CEO publicly estimated would require two to three years to repair. Beyond the energy shock itself, the conflict struck at the helium supply chain that underpins semiconductor manufacturing. Qatar produces roughly one-third of the world's helium, used to cool wafers during the etching process and to maintain inert environments in advanced lithography equipment. Damage to Ras Laffan caused output to fall approximately fourteen percent, with restoration estimated at up to five years. Major Asian chipmakers entered the conflict with three to six months of inventory buffer, according to JP Morgan, and helium remains less than one percent of production costs at current prices. However, a sustained Strait of Hormuz closure into the summer would translate dwindling reserves into a binding production constraint at a moment for an AI ecosystem already running on tight GPU supply.

Tungsten presents a parallel risk that has received insufficient attention relative to its severity. According to the US Geological Survey, China controls approximately seventy nine percent of global tungsten mine production and placed it on export control lists in February 2025. Prices have since risen more than 5.5x, with Chinese exports of restricted tungsten products falling roughly forty percent in 2025. Given tungsten's irreplaceable role in chipmaking at advanced nodes, appearing in both the chips themselves and the manufacturing equipment, analysts have forecast a supply deficit in 2026 with no near-term offset mechanism. The AI supply chain entered the conflict with almost no slack, as high-bandwidth memory was already sold out for the year and advanced chip packaging had backlogs extending 12 to 24 months.

The strategic positive from geopolitical stress is the intensified urgency of supply-chain diversification. The conflict opens a credible long-term investment case for semiconductor capacity in Europe and the US, a rebalancing that, over time, could reduce the systemic concentration risk currently making the AI buildout vulnerable to a handful of facilities in a narrow geographic corridor.

Watts per dollar: energy efficiency rewrites the competitive map

US structural supremacy and its limits: energy arbitrage, supply constraints and political risk

The US enters this environment from a position of structural advantage that the conflict has, paradoxically, reinforced. As the world's largest energy producer and net exporter, the US economy is relatively insulated from the supply shock. Henry Hub, the primary pricing point for US natural gas futures, reflects domestic supply and demand and has historically traded at a discount to global gas prices. According to the IEA, US prices for energy-intensive industries averaged less than half of those in the EU. Even before the Iran conflict, the IEA confirmed in its annual Electricity 2026 report that average EU wholesale electricity prices ran at roughly $95 per megawatt-hour, close to 2x US levels. Non-energy components, such as grid expansion fees, transmission network charges and local taxes, continue to make up a massive share of the final bill in Europe.

While the US and EU might have a wholesale price gap of roughly 1:2 ($48/MWh vs ~$95/MWh), the structural differences in grid tariffs and taxation mean the true gap in fully loaded industrial costs often scales differently across borders. Furthermore, according to a report by PowerLines, a consumer education group, US utilities are now planning $1.4 trillion in CapEx through 2030 to support AI-driven power demand growth.

Data centres currently consume 4.4 percent of US national electricity and the IEA projects a 130 percent increase from 2024 levels in US data centre consumption by 2030. The US is

building at scale, at competitive cost and with grid infrastructure that, while strained, has no equivalent in timeline flexibility outside of Southeast Asia. However, as reported by Bloomberg news, close to half of planned 2026 US data centre builds face delays from transformer and switchgear shortages, underscoring that even the most advantaged market is not immune to elongation risk. Nevertheless, the US retains the most favourable combination of energy cost, capital access, regulatory speed and foundational model leadership of any major economy.

US data centre buildout is scaling to match this energy advantage. Wood McKenzie anticipates capacity to grow from approximately 24 GW to 110 GW between 2026 and 2030, and accounting for 68% of total load growth over the period. Wood McKenzie estimates that, of the 600 GW of data centre projects still searching for power capacity in its US pipeline, only 183 GW have signed construction or electricity supply agreements with utilities.

US energy independence confers strategic moat status in the AI race. However, as outlined by PowerLines, it is a moat being paid for by American households. The midterm elections will likely turn on perceived cost-of-living and affordability issues and may influence the regulatory and rate-setting environment for the data centre buildout in ways that investors should not underestimate.

The EU's TTF tax on AI

Europe's AI ambitions are genuine; its implementation framework is not. The continent faces a structural energy cost problem that the Iran conflict has amplified. That problem is now intersecting with a dependency architecture that poses strategic risks of a different order.

The core arithmetic is unforgiving. According to the IEA's latest data, average electricity prices for energy-intensive industries in the EU hover around $107 per MWh, more than twice the US average of $50 per MWh, and roughly 50 percent above China. This structural disadvantage is exacerbated by the mechanics of Europe's liberalised electricity markets, which operate on a marginal pricing, or 'pay-as-clear,' model. Under this system, power is dispatched sequentially based on operating costs, starting with zero-marginal-cost renewables, moving to nuclear and ending with fossil fuels, until total demand is met. The critical mechanism is that the last, most expensive generation unit required to balance the grid sets the clearing price for all generators. Because flexible gas-fired plants frequently act as this final marginal unit, the European Title Transfer Facility (TTF) natural gas benchmark effectively dictates the wholesale cost of electricity. Consequently, when geopolitical shocks drive up TTF prices, industrial users absorb the full inflationary impact across the grid, while renewable and nuclear operators collect the high gas-cleared price without bearing the underlying fuel costs.

Grid interconnection lead times in established hubs stretch to ten years in some locations, according to Ember Energy. This renders most AI infrastructure planning exercises theoretical rather than operational. Across the world's major markets, average vacancy has compressed to just 6.3% by Q1 2026. In Europe's tier-one FLAP-D markets (Frankfurt, London, Amsterdam, Paris and Dublin), capacity is being absorbed significantly faster than it can be replaced. According to CBRE, the cost of securing future data centre capacity in these largest European hubs is projected to rise by twelve percent this year alone due to chronic supply constraints. Securing meaningful space now requires pre-commitment to construction pipelines rather than existing availability.

Two internal exceptions deserve note. The Nordic countries and France, beneficiaries of hydro, renewables and nuclear power respectively, offer meaningfully lower wholesale electricity prices than the EU average and are emerging as intra-European winners in the

competition for AI infrastructure. Norway has already attracted a $6.2 billion Microsoft - Nscale AI infrastructure deal. France's nuclear-anchored grid has led Google and Microsoft to direct significant European investment there. For equity investors seeking European AI infrastructure exposure, the geography of energy cost within Europe is a more important variable than the geography of Europe relative to the US.

By 2030, the IEA projects global data centre electricity consumption will more than double to reach approximately 945 TWh. Driven in large part by this infrastructure build-out, EU electricity demand is forecast to grow at an accelerated 2.3 percent annually through the end of the decade. If that growth relies on gas-fired generation, it compounds the inflationary pressure from the Iran shock and risks derailing Europe's climate transition simultaneously.

A Europe that cannot power its own AI infrastructure competitively, and that relies on American cloud for most of its computational needs, is not a sovereign AI actor, it is a client. The EU's permitting timelines, sometimes exceeding five years for data centre projects, and its fragmented grid infrastructure, which lacks the cross-border transmission capacity to equalise prices across the continent, mean that without structural policy reform, European AI ambitions will materialise only in the Nordics and France. The industrial heartland of Germany and central Europe may fall further behind.

Asia's paradox: convergence of manufacturing dominance and energy vulnerability

Asia occupies the most complex position among the three major AI regions. It is simultaneously the manufacturing core of the AI supply chain and one of its most exposed geopolitical and energy nodes. Taiwan's TSMC, South Korea's Samsung and SK Hynix and the advanced packaging ecosystem across Southeast Asia enables a substantial share of global AI infrastructure deployment. Yet the same region also carries disproportionate exposure to imported energy, semiconductor-material bottlenecks and geopolitical concentration risk.

Japan and South Korea illustrate this vulnerability most clearly. Both are structural energy importers with deep dependence on LNG and are exposed to the transmission of Middle East energy shocks through pricing, currency and supply-chain channels. In Japan, utilities rely heavily on LNG for power generation and price long-term supply contracts against the Japan Customs-Cleared crude benchmark. This creates a direct mechanical link to Brent crude with a three-to-six-month lag. Any weakness in the yen further amplifies this pressure, as LNG is priced in US dollars.

South Korea's exposure is more acute because it sits at the centre of global memory production. Samsung and SK Hynix, the world's two dominant memory-chip manufacturers, lost more than $200 billion in combined market capitalisation at the peak of the Iran conflict, as investors priced in both the direct energy-cost impact and the risk of helium disruption to semiconductor fabrication. South Korea imports roughly 70% of its crude oil from the Middle East, virtually all of it transiting the Strait of Hormuz.

The region's data centre market is therefore being reordered around power availability. CBRE's latest Asia Pacific analysis describes a market undergoing a significant shift, with direct data centre investment reaching a record $11.6 billion in 2025 and entity-level transactions expanding to $8.3 billion. As AI workloads accelerate, however, scalable power has become the decisive constraint. Capital is moving away from traditional Tier I hubs such as Singapore and Hong Kong, where capacity growth has slowed, and toward power-advantaged markets. Johor, Malaysia, for example, recorded a 53 percent y/o/y increase in live capacity in 2025, underscoring that grid access now often outweighs traditional location advantages for modern AI infrastructure providers.

The seven investable layers of the AI buildout

Having traced the build-out's capital intensity, its financing architecture and its energy and geopolitical fault lines, we close by translating the thesis into something investable. Each layer maps to a basket of publicly listed companies whose revenue profile, order book or capital allocation renders them efficient proxies for that specific segment of the build-out. These are constructed in a discretionary fashion, reflecting the view that thematic exposure is best captured by selecting the companies whose economics are most directly levered to each layer, rather than by relying on sector codes or revenue thresholds that would dilute the signal.

Each basket is market-capitalisation-weighted, giving the largest and most liquid representatives of each layer proportionate influence. Each is also deliberately concentrated: no fewer than three constituents and no more than five, selected for the purity of their exposure rather than the breadth of their coverage. For every constituent, we identify the segment or key driver that makes it an effective proxy and support that assessment with the latest financial disclosure available at the time of writing.

These seven layers connect sequentially and reinforce one another, as our May Equity Monthly Review argued. The fabrication layer determines the pace at which compute advances; the compute and memory layers define the economics of training and inference; the networking layer governs whether compute capacity can be utilised at cluster scale; the physical infrastructure layer translates silicon into deployable systems; the power layer determines the rate at which deployed systems can be energised; and the application layer determines whether the aggregate investment produces a commercial return. Investors seeking to construct thematic exposures are well-positioned to treat each of these layers as a discrete analytical unit, one that has its own revenue drivers, margin profiles and sensitivity to the pace of the build-out, rather than treating the AI complex as a monolithic trade. The baskets below are built to make that choice expressible.

The Fabrication Equipment & Process Materials Layer (L1 FE + PM)

The AI buildout does not begin at the chip; it begins at the tools and materials required to produce it. This layer encompasses the semiconductor equipment manufacturers and process-chemistry suppliers whose output is the necessary precondition for every compute advance downstream. Demand here is structural rather than cyclical, as leading-edge foundry-logic, high-bandwidth memory and advanced packaging are forecast to account for more than 80 percent of wafer fabrication equipment spending growth in 2026 alone.

This layer encompasses the companies that do not design chips, but sell the lithography, deposition, etch and process-control tools and the consumable materials, that every foundry and memory maker must buy. Its appeal is that it captures wafer volume regardless of which chip designer ultimately wins, and the cycle has validated the thesis that AI multiplies process complexity and consumable intensity rather than scaling them linearly.

- Applied Materials (AMAT) is the preeminent proxy for this layer. AMAT holds dominant positions in deposition, etch and metrology for gate-all-around transistors, DRAM and advanced packaging. Its Semiconductor Systems segment posted record revenue in Q1, with advanced-packaging revenues alone expected to grow more than 50 percent through 2026. The company's product introductions, including the Viva radical treatment system for nanosheet surfaces and the Kinex integrated die-to-wafer hybrid bonder, confirm that AMAT is not a passive beneficiary of AI CapEx, it is an enabling condition for the next node transition.

- ASML holds an uncontested monopoly on extreme ultraviolet lithography, the patterning technology required to manufacture leading-edge AI accelerators, high-bandwidth memory devices and sub-3-nanometre logic chips. There is no credible second source, no near-term substitute and no path to next-generation AI compute that does not run through Veldhoven. The company ended 2025 with a €38.8 billion backlog, nearly matching its full-year 2026 revenue guidance of €36 billion to €40 billion, confirming multi-year demand visibility independent of any single customer's capital-allocation cycle. To meet demand, ASML is increasing EUV tool shipments from 48 units in 2025 to at least 60 in 2026 and toward 80 by 2027, supported by a €2.2 billion CapEx programme to expand manufacturing capacity in Eindhoven. Beyond hardware, ASML's Installed Base Management (IBM) segment generated €8.193 billion in annual revenue in 2025, up 19.4% y/o/y, as component upgrades lift throughput without requiring new machine purchases. This creates a recurring-revenue stream that increasingly resembles a mission-critical services platform layered on top of a monopoly hardware base. Its next-generation High-NA EUV systems, priced at roughly $400 million each, will provide the patterning architecture for AI chip generations through the end of the decade.

- Lam Research offers the cleanest memory and high-bandwidth-memory leverage through its plasma etch and deposition equipment franchises, occupying two of the eight core processes required to manufacture a semiconductor. Its next-generation Akara etch platform, purpose-built for three-dimensional gate-all-around transistor architectures, and the ALTUS Halo atomic layer deposition tool, critical for void-free molybdenum metallisation at leading-edge nodes, position it at the precise inflection point where AI chip complexity translates into equipment intensity. Management estimates that that Lam's etch and deposition capabilities are disproportionately weighted toward advanced packaging and AI DRAM, the two fastest-growing subsegments of the cycle. Q1 revenue was a record $5.841 billion, up 23.8% y/o/y, Memory rose to 39 percent of its Systems revenue from 27 percent in Q4 and management lifted its 2026 Wafer-Fab-Equipment forecast to $140 billion.

- The materials dimension is held by Qnity Electronics, the pure-play electronic-materials business spun out of DuPont in November 2025 and added to the S&P 500, with Q1 sales of $1.315 billion, and Q1 EPS of $1.08. Qnity represents the consumables thesis in its most concentrated form. Its Semiconductor Technologies segment's revenue derives from recurring consumable products, chemicals, CMP pads, slurries and cleans, whose demand scales with wafer starts rather than with equipment capital cycles. Its Interconnect Solutions segment adds exposure to advanced packaging chemistry, thermal interface materials and high-layer-count printed circuit board solutions, precisely the substrate-layer products that AI's shift toward heterogeneous chiplet architectures is pulling into structural growth. Its collaboration with Nvidia on AI-driven materials modelling and simulation reinforces Qnity's role as a focused proxy for the materials-intensity theme within the AI buildout.

- Where equipment manufacturers supply the tools, Entegris supplies the consumable chemistry and handling systems that are consumed at every step of the fabrication process. Its three core growth drivers — CMP slurries and pads for inter-layer planarisation (Specialty Chemicals and Engineered Materials), ultra-high-purity chemical delivery and filtration systems (Microcontamination Control) and advanced wafer and reticle carriers (Advanced Materials Handling) — are not optional inputs; they are mandatory at each successive process step, and their per-wafer intensity increases with node complexity. As AI accelerators migrate toward three-dimensional DRAM architectures and heterogeneous advanced

packaging, the number of CMP and selective etch steps per wafer rises non-linearly. This creates a materials-intensity tailwind that the May Equity Monthly Review identified as one of the cycle's most defensible structural themes. Entegris' cross-licensing agreement with JSR's Inpria unit announced on 27 May on metal oxide resists for extreme ultraviolet lithography extends its addressable scope further into the patterning layer of next-generation chip fabrication.

Memory and Storage Layer (L2 M + S)

It is the tier that determines how quickly data can be presented to, held in proximity to and retrieved from accelerators, and it has emerged as one of the most structurally transformed segments of the AI buildout. High-bandwidth memory (HBM) has ceased to be a commodity input and become a designed, co-engineered component of the accelerator system itself: HBM content per GPU server rack increased 3.5x between the Hopper and Blackwell generations. The HBM market overall is projected to reach $54.6 billion in 2026, representing a 58% expansion over the prior year. Beneath the memory stack, the AI data pipeline generates an equally acute demand for enterprise flash storage, the medium through which training datasets are ingested, intermediate activations are checkpointed and inference outputs are logged at scale. Together, memory and storage constitute the bandwidth infrastructure without which the compute layer's raw processing power cannot be utilised.

As HBM consumes several times the wafer capacity of conventional memory, it is simultaneously the tightest supply bottleneck and one of the richest margin stories in the build-out. We separate it from compute because its pricing cycle and supply dynamics are distinct.

- Micron Technology is the cleanest US-listed expression. Its Q1 revenue reached a record $23.860 billion at a gross margin of 74.9%, with management guiding that figure toward approximately 81.0% in Q2 and projecting revenues of $33.500 billion. The company's HBM4 production capacity for calendar 2026 is sold out under long-term contracts. Its $200 billion domestic capacity investment commitment through the end of the decade positions it as the most consequential US infrastructure bet in the memory space. Critically, Micron's revenue sensitivity to AI is no longer expressed solely through HBM: AI servers consume substantially more DRAM and NAND than conventional servers, which means every incremental rack deployed in the buildout expands Micron's addressable opportunity across multiple product lines.

- SK Hynix is indispensable despite its Seoul listing. The company holds an estimated 50 percent or more of the global HBM market, a position Goldman Sachs expects to remain durable through at least 2026, and was the first manufacturer to ship HBM3E to NVIDIA for the Blackwell Ultra accelerator platform. Its dual-product strategy keeps HBM3E as the mainstream volume product while ramping HBM4 for next-generation deployments, giving it commercial exposure across both current and forthcoming GPU generations. SK Hynix was also featured at NVIDIA's GTC 2026 and HPE Discover 2026, where it displayed its HBM portfolio in direct physical integration with Vera Rubin platform hardware. That presence underscores that its relationship with the hyperscaler and accelerator ecosystem is not merely one of supply, but of co-design, creating a structural moat that commodity DRAM producers will find difficult to replicate. Q1 revenue reached a record KRW 52.576 trillion, up 198.1% y/o/y, and management has indicated that customer demand for the next three years already exceeds planned capacity, with a US depositary listing in progress.

- SanDisk, now an independent public company following its separation from Western Digital, represents the flash-storage layer of AI memory architecture. This layer supports data ingest, preparation, checkpointing and persistent storage across AI training and inference infrastructure. Its enterprise portfolio centres on the DC SN861 PCIe Gen 5 NVMe SSD, which delivers up to three times the performance of the previous generation, and is certified for NVIDIA MGX deployments. The UltraQLC platform, led by a 256 TB NVMe SSD built on BiCS8 QLC NAND, further strengthens SanDisk's position. The platform's direct-write architecture eliminates SLC caching, reduces write amplification and improves power efficiency by up to 10% through dynamic frequency scaling. SanDisk therefore acts as the storage-layer counterpart to Micron's memory-layer position: Micron supplies the bandwidth fabric adjacent to the accelerator, while SanDisk provides the high-density persistent storage that feeds it. Q1 revenue reached $5.950 billion, up 251.0% y/o/y, with data-centre revenue increasing 233.4% q/o/q and a multi-year contracted backlog exceeding $42 billion.

Compute and Silicon Layer (L3 C + S)

If fabrication is where chips are made, the compute layer is where chips are designed. It encompasses the GPU and XPU architectures. These perform the mathematical operations at the core of AI training and inference, including both merchant accelerator platforms and the growing ecosystem of custom application-specific integrated circuits commissioned by hyperscalers seeking performance-per-watt advantages unavailable from off-the-shelf configurations.

This is the layer where AI workloads are physically executed, and the demand origin for everything downstream: every accelerator installed pulls memory, networking, power and floor space behind it. It is also the most fully capitalised part of the chain, which is precisely why the marginal-investment debate has migrated lower down the stack.

- Nvidia is the irreducible anchor of this layer, whose Data Centre platform generated record revenue of $75.200 billion in Q1, up 92.3% y/o/y, with hyperscalers now accounting for only half of that figure and the balance drawn from sovereign, enterprise and AI-cloud demand, a broadening of the base that reduces customer concentration risk. A $95 billion supply backlog and the start of Vera Rubin sample shipments further reinforce that position.

- Advanced Micro Devices (AMD) provides the principal scaled alternative to Nvidia through its Instinct MI300 and MI400 accelerator families, supported by hyperscaler commitments to gigawatt-scale deployments across successive product generations through 2026. Its combination of merchant GPU supply and semi-custom system-on-chip co-design capabilities positions AMD both as a key beneficiary of any Nvidia supply constraint and as an active participant in hyperscaler compute-allocation decisions. The Data Centre segment generated $5.775 billion in Q1 revenue, up 57.2% y/o/y, while the agreement with Meta Platforms to deploy up to six gigawatts of Instinct accelerators anchors the MI450 and Helios rack-scale ramp into H2.

- Broadcom's role in the compute layer reflects the increasingly blurred boundary between merchant accelerator silicon and the custom Application-Specific Integrated Circuit (ASIC) programmes hyperscalers use to reduce reliance on GPU margins. The company holds an estimated 60 percent or more of the custom AI ASIC market, with confirmed programmes supporting Google's TPUs and Meta's MTIA. Its custom-silicon revenue offers hyperscalers a structural hedge as they optimise training and inference economics at scale inhouse, while its share of the silicon bill of materials in custom clusters makes Broadcom one of the compute

layer's highest value-capture companies. AI Semiconductor revenue reached $10.8 billion in Q1, up 143.0% y/o/y and equal to 48.6% of total revenue, supported by remaining performance obligations (RPO) of $164.6 billion and an AI order book above $30 billion.

- Marvell Technology's transformation from a networking chip supplier into a hyperscaler-aligned custom AI silicon platform is among the most consequential structural shifts in the sector. Its confirmed ASIC programmes with Amazon (Trainium), Microsoft (Maia), and Google (Axion) represent a revenue funnel management has guided toward $10 billion by fiscal 2029. The February 2026 acquisition of Celestial AI added photonic fabric optical interconnect technology, extending Marvell's addressable scope into the connectivity layer and creating a compound thesis that bridges the compute and networking layers. Marvell therefore completes the basket as a complementary custom-XPU and interconnect proxy, supported by $1.833 billion in data-centre revenue last quarter and a deepened NVIDIA partnership around NVLink Fusion, including a $2 billion investment.

Networking Layer (L4 N)

The networking layer is the connective tissue of the AI cluster. As GPU counts per training run scale from tens of thousands to hundreds of thousands of units, the latency, bandwidth and power efficiency of the interconnect fabric between accelerators become the binding constraints on system-level throughput. This layer encompasses the switching silicon, network interface cards, optical transceivers and end-to-end fabric architectures that ensure compute resources are never idle waiting for data.

- Arista Networks is the dominant provider of AI-optimised Ethernet switching platforms, including the 7060X6 AI leaf switch and the 7800R4 AI Spine switch, purpose-built for the scale-out requirements of large-scale GPU clusters. Its Etherlink AI platform has established it as the principal alternative to InfiniBand-based fabrics, and the company's end-to-end AI networking portfolio positions it to capture a disproportionate share of spending as open Ethernet architectures gain ground against proprietary interconnect solutions. Arista reported Q1 revenue of $2.709 billion, up 35.1% y/o/y, raised its full-year revenue outlook toward $11.5 billion, implying 27.7% growth, and lifted its AI-networking revenue target to $3.5 billion. It has also secured the number-one share position in high-speed switching above 10 gigabits per second (10GbE and higher).

- Cisco Systems represents the scaled incumbent being re-rated through AI infrastructure demand. The company secured $5.3 billion of hyperscaler AI-infrastructure orders over the past three quarters, including $1.9 billion in Q1 alone, and raised its full-year order target to $9 billion, roughly 4.5x the prior-year level. Its Acacia optics unit also booked more than $1 billion in orders in a single quarter, a record for this segment.

- Coherent is one of the market's most vertically integrated optical technology companies, spanning materials, devices, modules and systems across the full photonic stack. Its 1.6T and 3.2T transceivers and co-packaged optics solutions, showcased at OFC 2026 alongside a 400 Gbps-per-lane silicon photonics breakthrough with Tower Semiconductor, make it a key enabler of the bandwidth density required by next-generation AI switch ASICs. The company's vertical integration, from compound semiconductor substrate fabrication through to finished optical modules, provides a structural cost and performance advantage that is difficult for module-assembly competitors to replicate at scale. Its presence across both pluggable transceiver and co-packaged optics architectures ensures

relevance across successive generations of AI cluster design. Its Datacentre and Communications segment now accounts for 75.4% of revenue, which grew 17.3% y/o/y.

- Lumentum has become a leading Tier-1 supplier of indium phosphide-based laser components for AI data centre optics, with cloud and AI infrastructure now accounting for more than 60% of total company revenue. Its proprietary EML laser technology, which provides the light source for high-speed transceivers connecting GPU clusters at 800G and 1.6T data rates, remains capacity-constrained through 2026. This supply constraint has supported pricing discipline and accelerated the shift toward higher-margin 200G-per-lane EML products. The company's 1.6T DR4 OSFP transceiver prototype, unveiled at OFC 2026 alongside a new 1060 nm VCSEL-based optical interconnect platform, demonstrates a product roadmap aligned with the successive bandwidth generations hyperscaler AI infrastructure will require through the rest of the decade. In Q1, Lumentum recorded revenue of $808 million, up 90.1% y/o/y, with an optical-circuit-switch backlog already exceeding $400 million.

Site-readiness and Physical Infrastructure Layer (L5 SR + PI)

This layer captures the physical instantiation of the compute and networking layers, the server platforms, rack systems, liquid cooling architectures, and data centre integration capabilities that translate chip-level performance into deployable infrastructure.

It is the layer that the May Equity Monthly Review identified as the current binding constraint: Super Micro Computer's flagging of site-readiness as the bottleneck most likely to govern revenue recognition over the next two to four quarters is the definitive signal that physical infrastructure has migrated from commoditised input to critical path. An accelerator becomes productive only once it has thermal management, rack-level power distribution and a finished building to sit in.

- Vertiv sits at the precise intersection of the AI buildout's two most acute physical constraints: power delivery and thermal management. Its MegaMod HDX prefabricated infrastructure solution combines direct-to-chip liquid cooling with air-cooled architectures in modular deployments, supporting rack densities from 50 kW to more than 100 kW and total power capacity of up to 10 MW per module. This specification range is designed for GPU-dense AI factory environments. Beyond hardware, Vertiv's OneCore platform uses digital twin-based deployment orchestration to reduce on-site installation timelines by as much as 85% compared with traditional methods. In Q1, revenue rose 30.1% y/o/y to $2.650 billion, while the full-year revenue outlook was raised to $13.5 billion to $14.0 billion. Organic orders had increased 252% y/o/y in Q4 2025, driving total backlog to roughly $15 billion. The June 2026 acquisition of ThermoKey further strengthened Vertiv's heat-rejection portfolio and added heat-exchange technologies required to cool high-density AI data centres.

- Super Micro Computer has undergone a fundamental repositioning from traditional server manufacturer to full-stack AI infrastructure integrator. AI GPU-related platforms contributed more than 80% of revenues in Q1. The company's Data Centre Building Block Solutions portfolio now encompasses AI servers, storage, liquid cooling, power shelves and software management tools. Total revenue grew 122.7% y/o/y even as customer deployment delays linked to site readiness pushed a portion of recognised revenues into subsequent quarters.

- Equinix represents the real estate component of this layer as the world's largest data centre REIT by capacity and the leading provider of interconnection services.

Its footprint across 73 metros gives hyperscalers, sovereign AI programmes and enterprise customers the colocation density and carrier-neutral architecture needed to support distributed AI workloads. Its inclusion in the basket reflects the growing indispensability of neutral interconnection infrastructure as AI traffic expands across the data centre ecosystem. Roughly 60% of Equinix's largest deals are now AI-related, driven by 8 of the top 10 AI model providers and 4 of the top 5 neo-clouds expanding inside its facilities. Fabric interconnection revenue grew 26.0% across more than 513,000 connections, supported by a 70% increase in related bookings as customers scaled network capacity for distributed AI workloads.

- Digital Realty provides large-scale data centre campuses across the Americas, EMEA and Asia-Pacific, with a development pipeline in the Americas alone representing hundreds of megawatts of announced future capacity as of Q1. Its positioning as a wholesale data centre provider to hyperscalers and AI-native cloud companies makes it a direct real-asset proxy for the physical footprint of the build-out. While the REIT structure introduces rate sensitivity, the structural demand signal from the AI capex cycle provides a durable occupancy and pricing floor unavailable to generic real estate platforms. Digital Realty signed its largest-ever lease, a 200-megawatt hyperscale agreement in Charlotte, North Carolina, lifting backlog to a record $1.8 billion against a development pipeline of 1.2 gigawatts.

Power and Electrification Layer (L6 P + E)

Power has graduated from a financing footnote to the primary constraint on the pace of the AI build-out. GE Vernova's Q1 disclosure that data-centre-related electrification orders exceeded the entirety of 2025 in a single quarter, combined with Constellation Energy's repositioning as a structured counterparty to the CapEx cycle, establishes that clean baseload generation and grid delivery are the binding variable determining when committed CapEx converts to operational infrastructure.

We identified power and site-readiness as the marginal constraints on the pace at which CapEx converts to revenue. The order data now bears this out emphatically. The strongest book-to-bill and backlog signals in the chain are concentrated here, not in silicon. This layer covers both generation — the electrons themselves — and the electrical equipment that moves power from the grid to the rack.

- GE Vernova is the most direct public-market proxy for AI-driven power infrastructure demand. Its gas turbine backlog increased from 62 GW to 83 GW across 2025 and early 2026, while data-centre-related electrification orders reached $2.4 billion in Q1 alone, exceeding the $2.2 billion booked during all of 2025. The company recorded total Q1 orders of $18.3 billion, representing a 71.0% organic increase y/o/y and implying a book-to-bill ratio of 1.97x, and ended the quarter with a $163 billion backlog. GE Vernova monetises the AI buildout through the equipment required to deliver power to data centres, including turbines, grid solutions, substations and energy management systems. Its first energy management system order for a data-centre customer, closed in Q1, marks a meaningful expansion of its addressable market.

- Eaton sits at the critical interface between utility-scale power delivery and the internal electrical architecture of the data centre. Its portfolio spans medium- and low-voltage switchgear, large-frame uninterruptible power supply systems, busway and prefabricated modular electrical rooms. It covers the full chain from utility interconnection to rack-level power distribution. Data centre-related orders rose approximately 240% y/o/y through 2025 and into 2026, creating a 228 GW

backlog signal that provides substantial multi-year revenue visibility independent of any single hyperscaler's quarterly CapEx cycle. The January 2026 collaboration with Flexnode to deploy Eaton's 800 VDC power systems in prefabricated AI-focused data halls, shortening deployment schedules by roughly 35%, together with a $30 million investment in a new medium-voltage switchgear facility near Omaha, Nebraska, underscores a capital-allocation programme aligned with the pace of the AI buildout rather than the volatility of commodity markets.

- The generation dimension is held by Constellation Energy. The company has emerged as the canonical clean baseload counterparty for hyperscaler power procurement. Its nuclear generation fleet provides the firm, carbon-free power profile that AI operators require for sustainability commitments and grid reliability. The February 2026 agreement between its Calpine unit and CyrusOne, covering more than 1,100 MW to support a new data centre facility at the Freestone Energy Center in Texas, exemplifies the structured, long-duration power purchase agreements that differentiate Constellation Energy from merchant generators exposed to spot pricing. Constellation Energy's nuclear fleet achieved a 92.3% capacity factor for Q1, generating 40 million megawatt-hours of zero-carbon electricity. The structural thesis remains intact as the cleanest large-scale baseload option available to the hyperscaler procurement cycle.

- Vistra operates the largest competitive power-generation fleet in the US, combining gas-fired assets with an expanding nuclear portfolio following its acquisition of Energy Harbor. Its presence across deregulated markets in Texas, the Mid-Atlantic, and the Midwest places it in the regions where AI data centre load growth, broader electrification demand and grid congestion are tightening supply-demand balances most sharply. Vistra complements GE Vernova's infrastructure-equipment thesis by providing direct generation exposure. GE Vernova builds the hardware that enables power delivery, while Vistra sells the electrons that flow through it. The company has signed roughly 3,800 megawatts of nuclear power-purchase agreements with Amazon Web Services, as well as multi-gigawatt, twenty-year agreements with Meta Platforms.

Application Layer (L7 A)

The preceding six layers describe the supply of the build-out; the seventh is where it must ultimately be paid for. The application layer is the demand signal that underwrites the entire chain. The May Equity Monthly Review flagged its first credible monetisation as the cycle's pivotal unknown. It is also the layer whose monetisation has been most interrogated by the investment community, and whose validation, still early but increasingly credible, may prove to be the next significant re-rating catalyst for the broader AI complex. The early read from this layer's constituents is that the conversion from pilot deployment to contracted recurring revenue is proceeding faster than consensus had modelled.

- Salesforce's Agentforce platform reached $800 million in annual recurring revenue (ARR) in Q1, representing the first commercially validated evidence that an agentic AI product can achieve scale within an enterprise software context. The introduction of the Agentic Enterprise License Agreement, a flat-fee, outcome-aligned pricing structure designed for organisations deploying AI agents at scale, signals a deliberate transition from per-seat SaaS economics. It is moving towards a model in which revenue scales with the work performed rather than the headcount supported.

- ServiceNow's Now Assist surpassed $600 million in Annual Contract Value (ACV) by end-2025. It is on track to approach $1 billion by the end of 2026, making it the fastest product ramp in the company's history. Its shift from per-seat licensing to

a hybrid model, combining a guaranteed revenue floor with consumption-based components tied to assets, devices and data volumes, offers a template for enterprise AI monetisation that is less dependent on headcount and more closely aligned with operational throughput. With full-year 2026 subscription revenue guidance nearing $15.755 billion, ServiceNow's AI-native product layer has become a material and expanding part of its compounding revenue base. In Q1, subscription revenue reached $3.671 billion, up 22.1% y/o/y, with remaining performance obligations (RPO) of $27.700 billion and the number of Now Assist customers spending more than $1 million in ACV growing by more than 130% y/o/y.

- Microsoft occupies a dual role in the AI taxonomy, straddling the cloud infrastructure layer as Azure's operator and the application layer as the most broadly deployed enterprise AI platform through Copilot. Its inclusion here reflects the application layer's logic: the commercial conversion of AI infrastructure into billable enterprise outcomes is, at present, disproportionately channelled through Microsoft's productivity, CRM and developer tooling surface area. Azure OpenAI Service, GitHub Copilot and Microsoft 365 Copilot collectively provide the widest enterprise distribution network of any AI application suite. The $190 billion CapEx programme for 2026 simultaneously reinforces the infrastructure thesis and funds the application-layer buildout from a single balance sheet. The signal to watch is whether application-layer revenue continues compounding at triple-digit rates: sustained growth validates the CapEx flowing into every layer beneath it, while a stall would re-rate the most capital-intensive layers first.

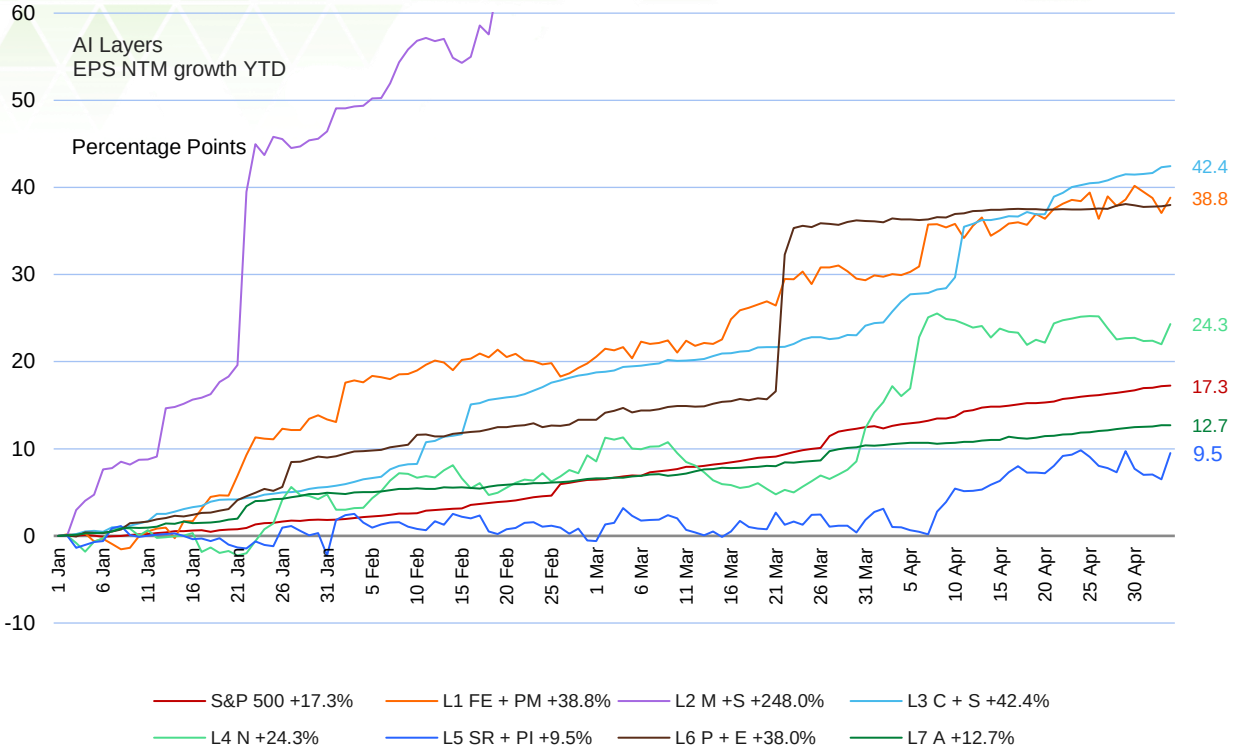
Two caveats apply across the framework. Several constituents, Broadcom, Cisco, Eaton, Constellation and Microsoft among them, are not pure-plays and are included because their AI-exposed segments are the marginal earnings drivers rather than the majority of revenue; their thematic sensitivity is therefore lower than that of the pure-play names. This framework ultimately relies on sustained hyperscaler CapEx; any deceleration in webscale outlays would trigger a simultaneous contraction in order volumes and valuation multiples across the seven layers, positioning the application layer as a critical leading indicator for the downstream infrastructure.

How is each layer positioned in the market?

Growth trajectory and profitability: which layers are earning their place

The four exhibits that follow read together rather than in isolation. The first traces how forward earnings expectations have re-rated since the turn of the year. The next two decompose profitability into its EBITDA and free-cash-flow components and, critically, their direction of travel; and the scatter pairs forward EV/EBITDA against expected EBITDA growth to ask the only question that ultimately matters for an allocator, whether the price being paid is warranted by the growth being bought.

The headline is that the market has not re-rated the build-out uniformly. Consensus next-twelve-month (NTM) earnings for the seven baskets have moved from a common base on 31 December to a remarkable dispersion by late June: Memory and Storage (L2 M +S) has seen forward earnings revised up by 248.0%, Compute and Silicon (L3 C + S) by 42.4%, Fabrication (L1 FE + PM) and Power by 38.8% and 38.0% respectively, Networking (L4 N) by 24.3%, the Application layer (L7 A) by 12.7% and Site-readiness (L5 SR + PI) by 9.5%. This is not a thematic tide lifting all boats equally; it is a sharp reordering of where the earnings power of the build-out is actually accruing.



Profitability direction sharpens the picture. The exhibits confirm that the layers closest to the silicon are not merely growing but becoming structurally more profitable: Compute (L3 C + S) expands its EBITDA margin to 58.0% and Memory (L2 M + S) to 65.1% at the basket level, the latter a swing of nearly 16 percentage points y/o/y as high-bandwidth memory pricing converts directly into operating leverage. Fabrication (L1 FE + PM) holds margins in the mid-thirties with FCF intact, consistent with its description earlier as the layer that captures wafer volume regardless of which designer wins. The picture inverts as one descends toward the physical layers. Site-readiness (L5 SR + PI) shows margins compressing and aggregate FCF conversion turning negative, and Power's (L6 P + E) reported basket margin contracts even as orders surge — a tension explored below. Investors should treat the layer-level margin aggregates with care, as they are market-cap weighted averages, and in the thinner baskets, a single constituent distorts the headline. Super Micro's low-single-digit margins drag the Site-readiness average downward, and the absence of complete REIT data for Equinix and Digital Realty, whose economics are properly measured on FFO rather than EBITDA, and which are not yet captured here, means that basket is best read through its constituents rather than its average.

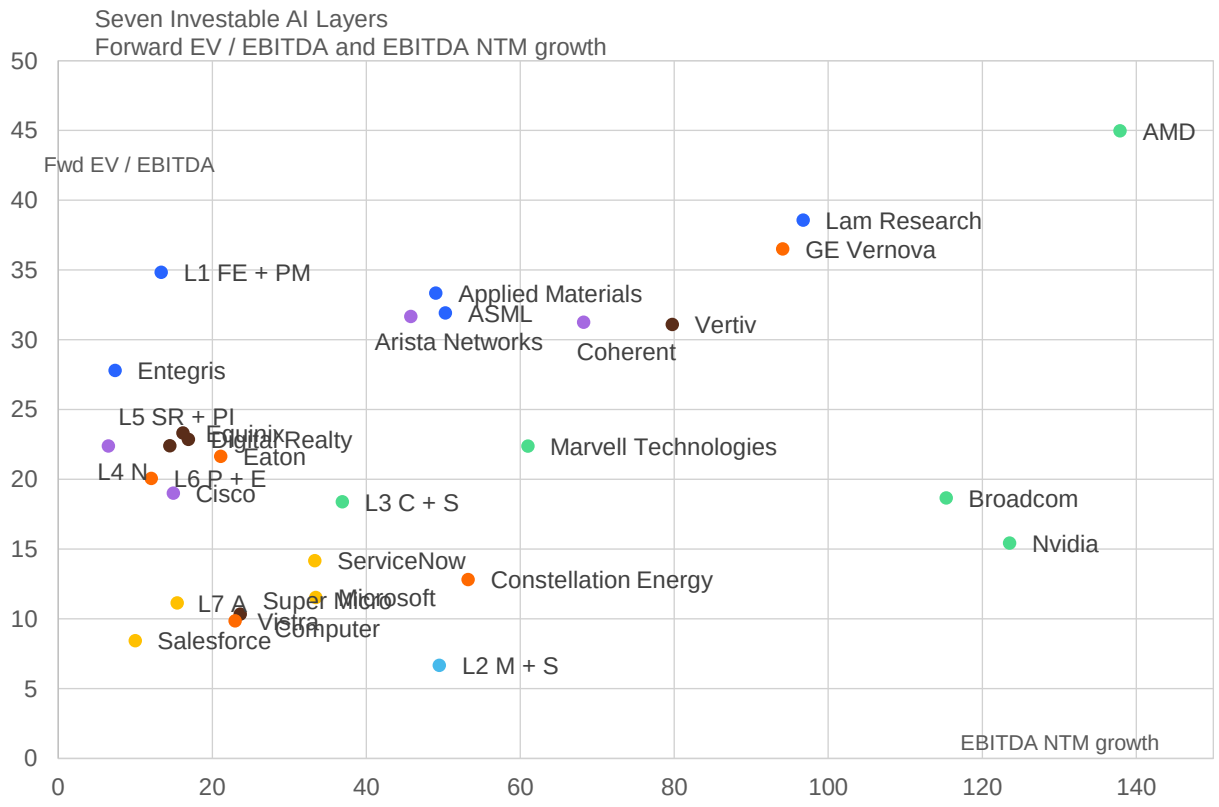
Taken together, growth and profitability identify where the build-out is genuinely earning its place. Compute and Memory are growing fastest and becoming more profitable simultaneously, the most favourable combination on offer. Fabrication grows at a respectable pace while defending high, stable margins, the profile of a durable toll-collector. The Application layer, though the slowest to re-rate, is the only tier where the profitability of monetisation — not the promise of it — is now visible in expanding software margins. The physical layers, by contrast, are being asked to invest ahead of return: their order books are full, but the cash conversion that would validate the spending has not yet arrived.

Valuation and the price of growth: are only blue skies priced in?

The scatter of forward EV/EBITDA against expected EBITDA growth is where opportunity and risk become legible, and it cautions against the lazy conclusion that the whole complex is expensive.

The scatter arranges each layer and its constituents along two axes: expected EBITDA growth on the horizontal, and the forward EV/EBITDA multiple the market is currently paying on the vertical. Read together, they describe the price of growth. A name in the lower-right quadrant, rapid expected growth at a modest multiple, offers the most favourable terms, while a name in the upper-left, where an elevated multiple accompanies more pedestrian growth, leaves the least room for error. The instructive cases are those that defy the expected upward-sloping relationship. Memory and Storage (L2 M + S) stand out in the lower-right quadrant: triple-digit EBITDA growth paired with a single-digit multiple, reflecting the market's view that this growth is potentially cyclical rather than durable. The Compute anchors occupy the opposite logic, pairing genuinely high growth with multiples that remain reasonable for the franchises concerned, whereas the premium fabrication, networking and power names carry multiples that already discount sustained execution. The diagonal, in effect, separates the layers whose growth is in the price from those whose growth is not yet fully paid for.

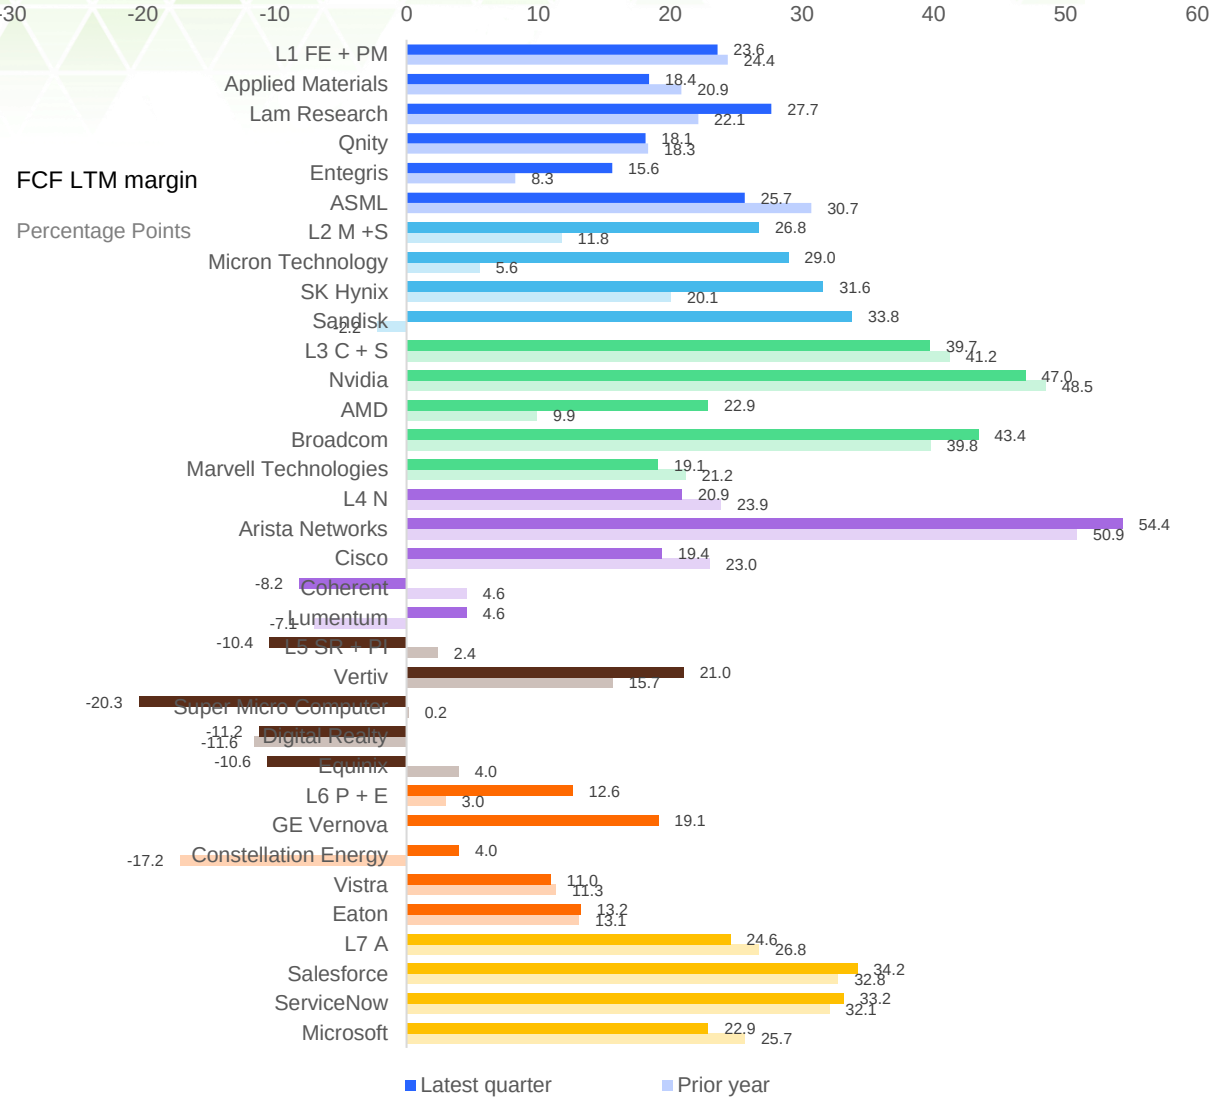
The standout is Memory and Storage, whose position as an outlier is entirely explained by its constituents. Micron, SK Hynix and Sandisk are lapping the trough of a brutal memory down-cycle, so their NTM EBITDA growth screens at several hundred percent (Micron above 600 percent, Sandisk above 2,000 percent) while the multiple the market will pay sits at a strikingly low five to nine times. The result is a layer that has re-rated its forward earnings by 248.0%, yet trades on a single-digit EV/EBITDA, the cheapest growth in the entire build-out on a mechanical reading. The driver is unambiguous: the transition of high-bandwidth memory from a commodity input to a co-engineered component of the accelerator system itself has compressed the supply-demand balance to a degree that has allowed Micron and SK Hynix to reprice entire product lines simultaneously. The caveat, however, is that the growth implied might be cyclical, not secular.

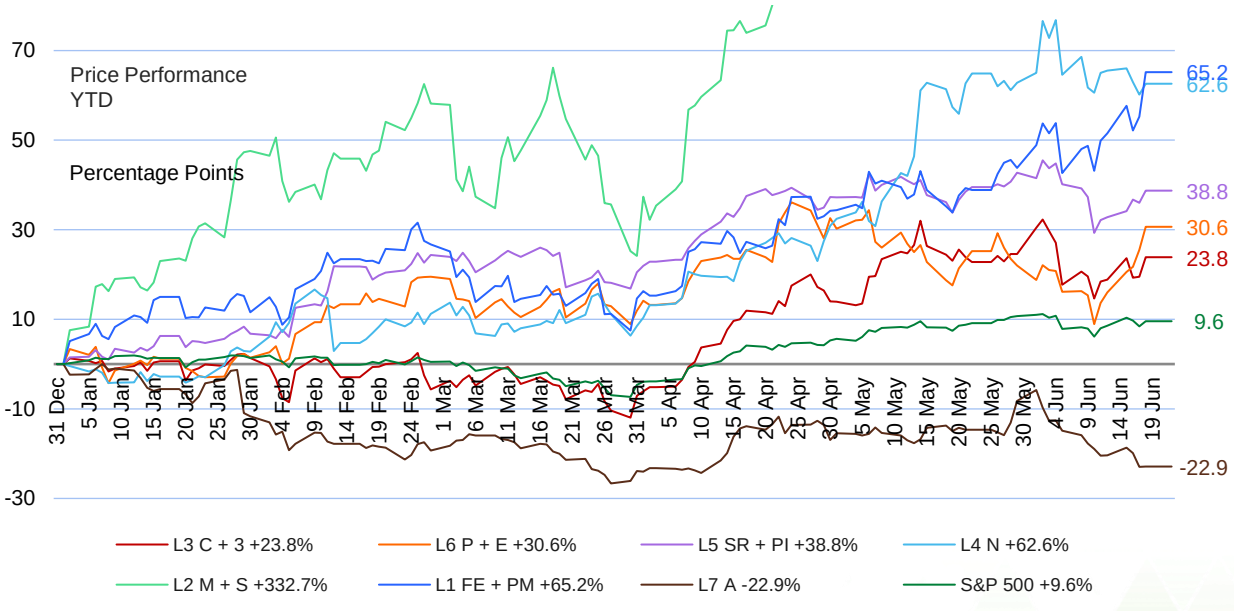


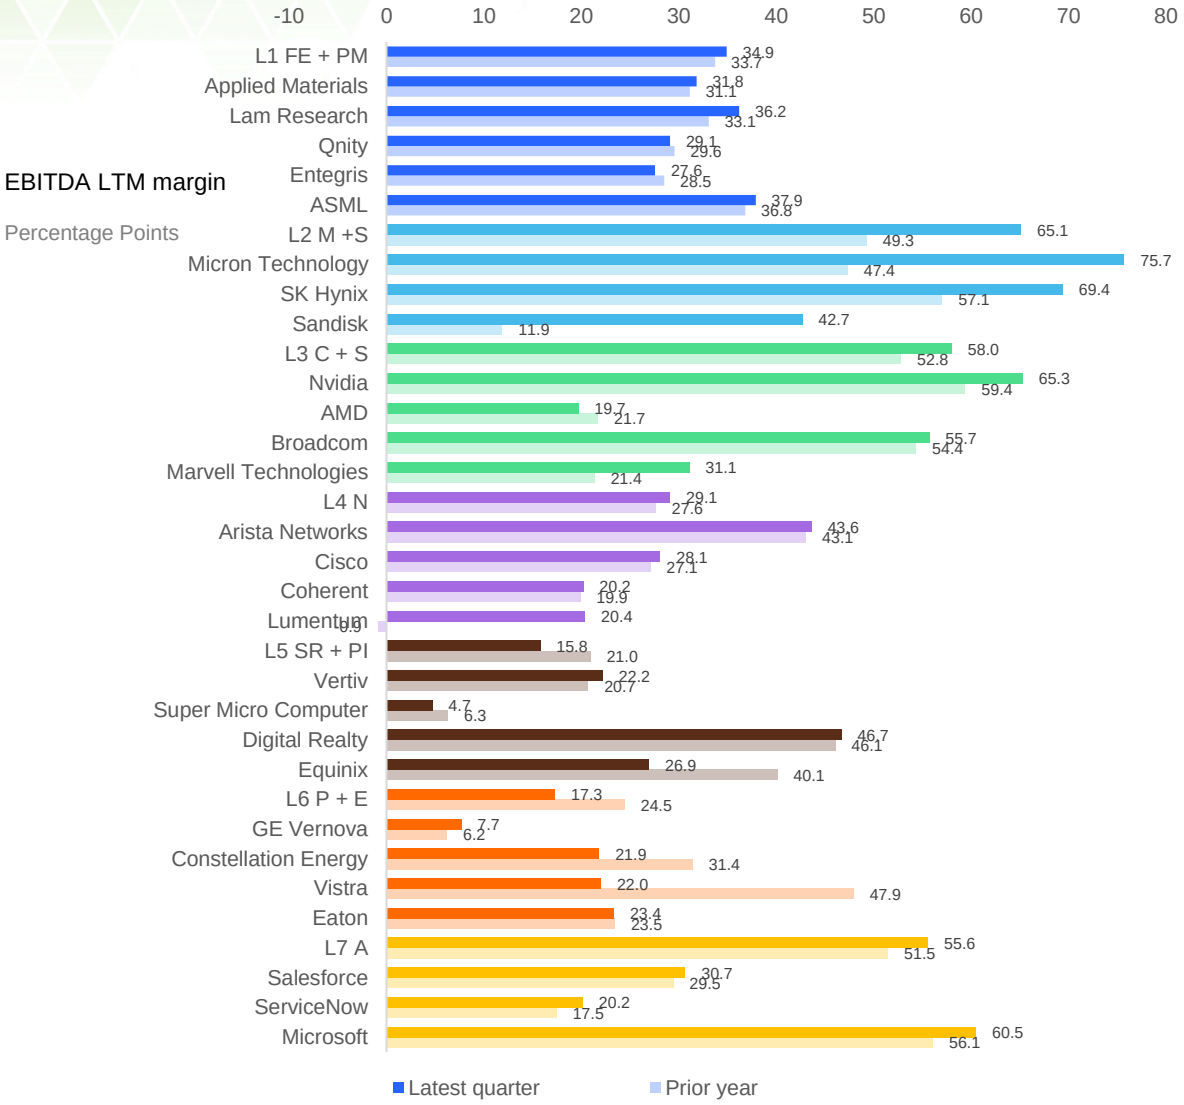
Elsewhere the scatter rewards discrimination over conviction. Compute and Silicon (L3 C + S) offers the most defensible combination of scale, growth and price, as Nvidia anchors the layer on a mid-teen forward multiple despite triple-digit growth, an unusually reasonable

valuation for the franchise at the centre of the theme, while Broadcom's custom-silicon economics screen similarly. The premium names within the layer, AMD and Marvell at 40x to 46x, carry the valuations where execution must be flawless, and where disappointment would be punished first. The same dispersion governs the other layers. Within Networking (L4 N), Cisco re-rates modestly at nineteen times while Arista and the optical names command low-thirties multiples that already discount sustained AI-fabric share gains. Within Power (L6 P + E), GE Vernova trades at thirty-six times against a high-twenties to mid-thirties multiple for the layer, pricing the order-book momentum described earlier rather fully. The independent power producers, Vistra and Constellation, remain in single-digit to low-teens territory despite signing the same hyperscaler load. Vistra's apparent margin contraction is a reporting artefact of hedging and acquisition accounting rather than deterioration. Constellation's improving FCF conversion is masked at the basket level.

The synthesis for a forward-looking investor is that blue skies are not uniformly priced in, but they are priced in selectively; the distinction is where the opportunity lies. Genuine value sits where durable, secular growth has not been fully capitalised: the toll-collectors of Fabrication, and the still-cheap independent power producers attached to multi-decade demand. The clearest risk sits where elevated multiples already embed flawless execution, the premium compute names, the optical networking specialists and where any deceleration in hyperscaler CapEx would compress both the multiple and the order book at once. Memory occupies a category of its own: spectacular on every backward- and near-forward-looking measure. As argued from the outset, the question is less whether the build-out is real than which layer to own, and at what price. Although the exhibits do not answer that question on investors' behalf, they aim to frame it with unusual clarity.


FCF LTM margin
Percentage Points
-30
-20
-10
0
10
20
30
40
50
60
L1 FE + PM
Applied Materials
Lam Research
Qnity
Entegris
ASML
L2 M +S
Micron Technology
SK Hynix
Sandisk
L3 C + S
Nvidia
AMD
Broadcom
Marvell Technologies
L4 N
Arista Networks
Cisco
Coherent
Lumentum
L5 SR + PI
Vertiv
Super Micro Computer
Digital Realty
Equinix
L6 P + E
GE Vernova
Constellation Energy
Vistra
Eaton
L7 A
Salesforce
ServiceNow
Microsoft
23.6
24.4
18.4
20.9
27.7
22.1
18.1
18.3
15.6
8.3
25.7
30.7
26.8
11.8
29.0
5.6
31.6
20.1
33.8
2.2
39.7
41.2
47.0
48.5
22.9
9.9
43.4
39.8
19.1
21.2
20.9
23.9
54.4
50.9
19.4
23.0
-8.2
4.6
-7.1
4.6
-10.4
2.4
21.0
15.7
-20.3
0.2
-11.2
-11.6
-10.6
4.0
12.6
3.0
19.1
-17.2
4.0
11.0
11.3
13.2
13.1
24.6
26.8
34.2
32.8
33.2
32.1
22.9
25.7
Latest quarter
Prior year


Price Performance
YTD
Percentage Points
70
50
30
10
-10
-30
65.2
62.6
38.8
30.6
23.8
9.6
-22.9
31 Dec
5 Jan
10 Jan
15 Jan
20 Jan
25 Jan
30 Jan
4 Feb
9 Feb
14 Feb
19 Feb
24 Feb
1 Mar
6 Mar
11 Mar
16 Mar
21 Mar
26 Mar
31 Mar
5 Apr
10 Apr
15 Apr
20 Apr
25 Apr
30 Apr
5 May
10 May
15 May
20 May
25 May
30 May
4 Jun
9 Jun
14 Jun
19 Jun
L3 C + 3 +23.8%
L6 P + E +30.6%
L5 SR + PI +38.8%
L4 N +62.6%
L2 M + S +332.7%
L1 FE + PM +65.2%
L7 A -22.9%
S&P 500 +9.6%



Horacio Coutiño
Multi-Asset Strategist
hc@exante.eu

DISCLAIMER: This article is provided to you for informational purposes only and should not be regarded as an offer or solicitation of an offer to buy or sell any investments or related services that may be referenced here.

While every effort has been made to verify the accuracy of this information, EXT Ltd. (hereafter known as "EXANTE") cannot accept any responsibility or liability for reliance by any person on this publication or any of the information, opinions, or conclusions contained in this publication. The findings and views expressed in this publication do not necessarily reflect the views of EXANTE. Any action taken upon the information contained in this publication is strictly at your own risk. EXANTE will not be liable for any loss or damage in connection with this publication.