

MEMO NO. 2 — ON THEMATIC INTELLIGENCE

The Largest Infrastructure Buildout in Human History.

Has Your Advisor Mentioned It?

McKinsey and JPMorgan project cumulative AI infrastructure investment of \$5–7 trillion by the end of the decade. The bottlenecks are not in software — they are in energy, chip packaging, memory fabrication, and the physics of moving data.

At the World Economic Forum in January 2026, Jensen Huang — founder of the world's most valuable semiconductor company — described artificial intelligence not as a technology but as a five-layer industrial stack: energy at the base, then chips, infrastructure, models, and applications at the top. Every successful application, he said, pulls on every layer beneath it, all the way down to the power plant.

This was not a product pitch. It was a description of a capital expenditure cycle without modern precedent. The five largest hyperscalers alone plan to spend approximately \$690 billion in 2026 — nearly tripling their 2024 outlay. Add the upstream supply chain — semiconductor fabrication, memory manufacturing, energy infrastructure — and the annual figure is substantially larger. McKinsey and JPMorgan project cumulative AI infrastructure investment of \$5–7 trillion by the end of the decade.

This matters to anyone with a portfolio. Not because you should pick stocks in any of these layers, but because this buildout is restructuring the economy beneath your existing holdings in ways that most advisory chains are not equipped to see — let alone position for.

"AI is infrastructure. Every country should treat it like electricity or roads."

— Jensen Huang, WEF Davos, January 2026

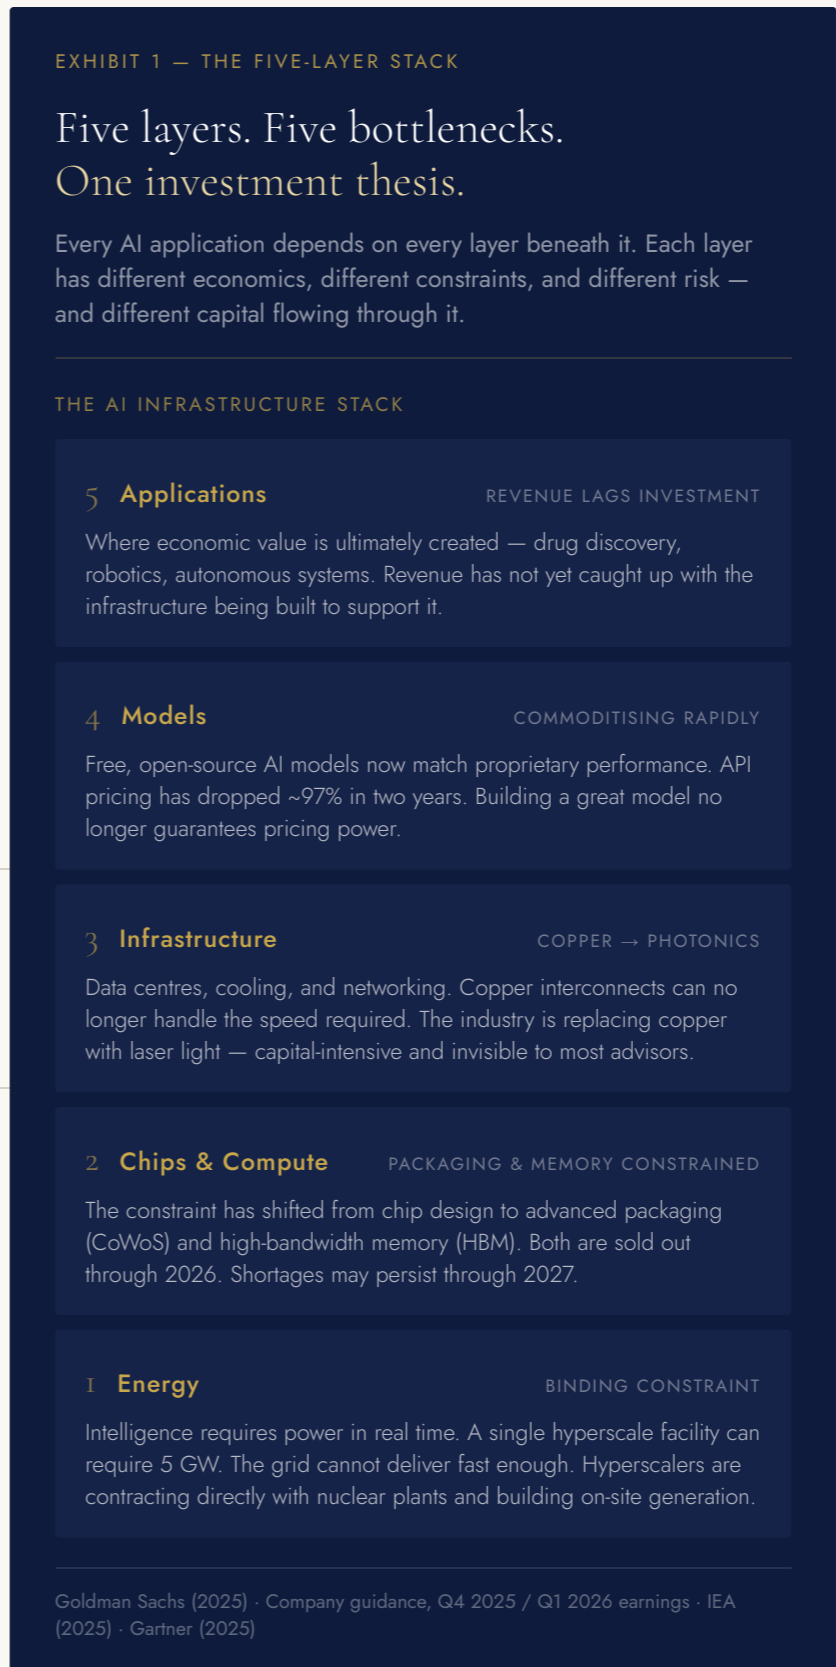
WHY THIS IS NOT A TECH THEME

Most banks sell "AI exposure" the way they sold "internet exposure" in 1999 — through a basket of familiar names, repackaged under a thematic label. AI is not a sector. It is an industrial chain with five distinct layers, each with different economics, different risks, and different bottlenecks. Owning a software company is not the same as owning the energy infrastructure that makes the software possible.

THE ENERGY CONSTRAINT

Intelligence generated in real time requires power generated in real time. There is no abstraction layer beneath this. US data centre power demand currently sits below 15 gigawatts. The pipeline of planned facilities would more than triple that figure. Gartner projects global data centre electricity consumption will rise from 448 terawatt-hours in 2025 to 980 by 2030. Hyperscalers are signing direct agreements with nuclear plants, restarting decommissioned reactors, and building behind-the-meter generation because the grid cannot deliver fast enough.

A portfolio with "AI exposure" and no view on the energy layer is incomplete.



THE CHIP LAYER AND WHAT SITS BENEATH IT

The dominant GPU manufacturer captures approximately 90% of AI accelerator spend. That is a remarkable franchise — and a remarkable concentration risk. Beneath it sits a supply chain of memory, advanced packaging, optical interconnects, and photonics that determines how fast data moves between processors. This is where some of the most consequential bottlenecks sit. Most thematic products never get this deep.

Advanced packaging — specifically TSMC's CoWoS process — is the physical step that connects logic chips to high-bandwidth memory. HBM itself consumes three to four times the silicon wafer area of conventional memory but generates more than ten times the revenue per wafer. Both are sold out through 2026. The companies that fabricate this memory represent less than 0.2% of global semiconductor wafer starts but capture roughly 20% of industry revenue. That asymmetry is invisible to most advisory chains.

INFRASTRUCTURE IS NOT CLOUD

An AI data centre is not a server farm. It is a factory for producing intelligence — requiring liquid cooling at industrial scale, power delivery systems designed for constant load, and networking architecture that connects tens of thousands of processors into a single machine. A single hyperscale facility can cost more than \$10 billion. These are physical assets with construction timelines, permitting risk, and energy contracts that resemble industrial infrastructure more than technology companies.

The critical transition underway: data inside these facilities currently moves through copper wires. But copper interconnects are hitting a physical ceiling — at 224 gigabits per second, signals can only travel about one metre before degrading. The replacement is silicon photonics: laser light through glass fibre, capable of vastly higher speed and consuming roughly 30% less energy. In March 2026 alone, over \$4 billion was invested in silicon photonics companies. When your bank says "cloud," ask whether they mean software or concrete.

LAYER 4 — MODELS: THE PRICE OF FREE

In January 2025, a Chinese research lab released an open-source AI model that matched frontier performance at a reported training cost of \$5.6 million — a fraction of what Western labs spend. Markets erased \$600 billion in technology market capitalisation in a single session. Since then, the gap between free, open-source models and the most expensive proprietary ones has narrowed to under 2%. API pricing has fallen approximately 97% since 2022.

The implication: if your AI exposure depends on a company maintaining pricing power at the model layer, the structural economics are working against it. Value is migrating toward the infrastructure that runs the models and the applications that deploy them.



LAYER 5 — APPLICATIONS: CUSTOMERS, NOT WINNERS

Drug discovery platforms. Industrial robotics. Autonomous vehicles. Legal and financial copilots. Software that writes software. The application layer is where economic value will ultimately be created — and where the \$690 billion in annual infrastructure investment must eventually find its return.

It has not yet. Several hyperscalers now spend more on AI infrastructure than they generate in free cash flow. They are funding the gap with debt issuance — transforming historically cash-funded businesses into leveraged ones for the first time. The bet is that AI-powered services will produce trillions in revenue over the next decade. That may prove correct. But it is a forward-looking bet, not a settled fact.

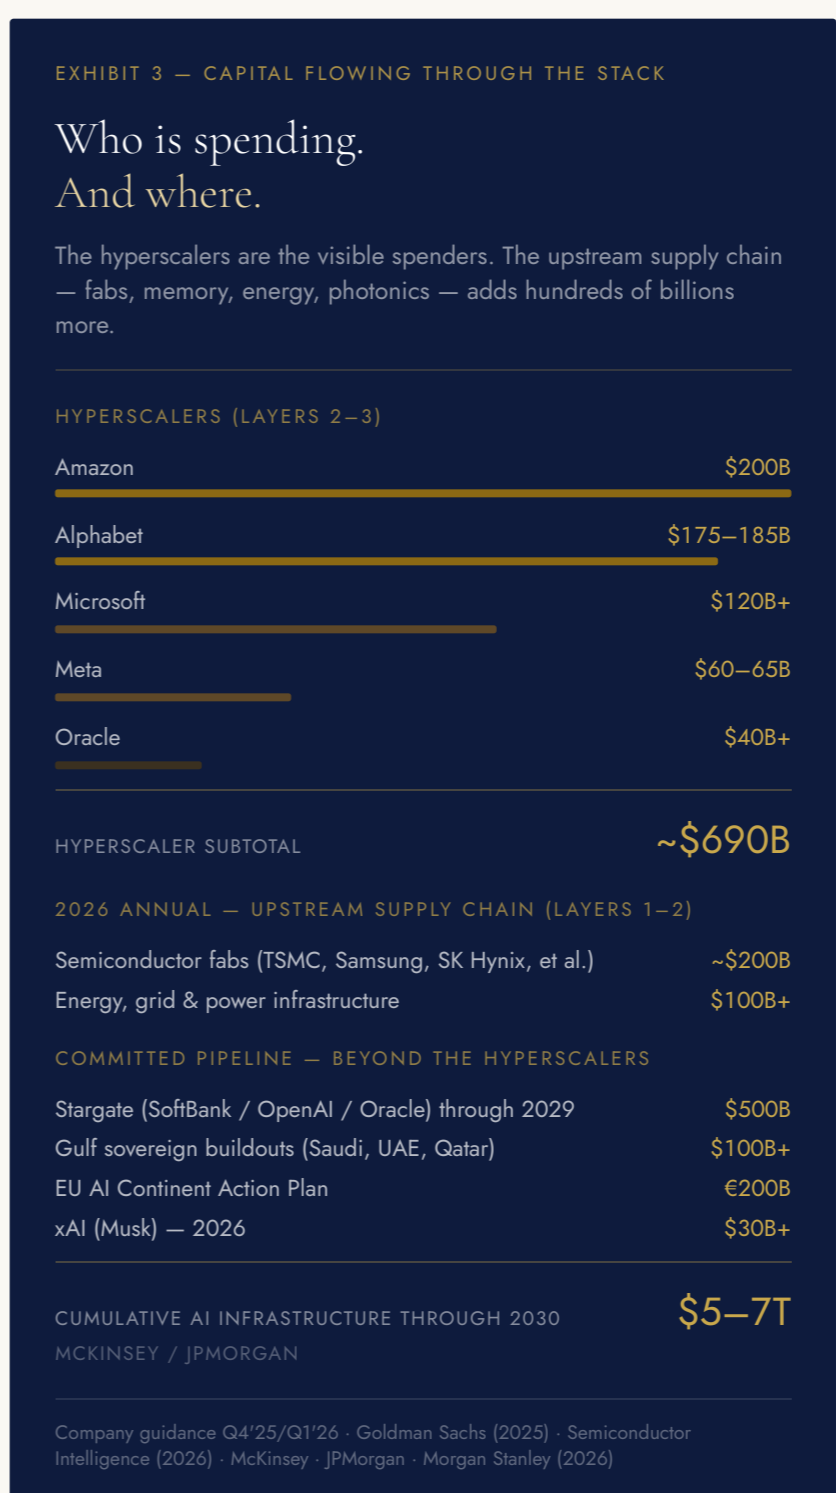
Here is the distinction most thematic positioning misses: the five companies spending \$690 billion are not the AI suppliers. They are the AI customers. They must buy the energy, the memory, the chips, the packaging, and the photonics from companies that most advisors have never heard of — and those suppliers are sold out, raising prices, and generating real profits today. Not in a forward model. Now.

Buying the five largest hyperscalers as "AI exposure" is like buying JPMorgan in 2015 as a "fintech play" — because it was the largest bank and the safest-sounding name adjacent to the theme. But JPMorgan was not fintech. It was the incumbent that fintech was aimed at. The same structural question applies: are you owning the theme, or the incumbents who are spending to survive it?

THE GOVERNANCE QUESTION

This is not a recommendation to buy or sell anything. It is an observation about governance: the largest capital reallocation in a generation is underway, and most private wealth portfolios are positioned for it using the same instruments and the same analysis that worked in 2015. The five-layer framework is a diagnostic tool. If your advisor cannot explain where your portfolio sits across all five layers — and where it is exposed, concentrated, or absent — then you have a governance gap, not a product gap.

The solution is not more products. It is a better question. And the question is not *which AI stocks should I own*. The question is: *does anyone in my advisory chain actually understand the architecture of what they are putting me into?*



THE CIOffice PROPOSITION

The Intelligence Layer. Architecture Alpha.

This memo demonstrates a structural analysis that sits above any individual advisor, fund, or product — and that most advisory chains are not built to deliver.

Consider a concrete example. Your bank puts you into a thematic AI fund. It holds the same ten mega-cap names that already dominate your core portfolio. Nobody noticed, because nobody consolidated. That is not just a cost problem. That is a performance problem. You are paying for diversification you do not have — and you are absent from entire layers of the buildout where real scarcity creates real pricing power.

CIOffice does not replace your advisors. We do not pick stocks. We provide the intelligence layer — the consolidated architectural view that makes every advisor across your structure more effective. *We don't just help you stop losing. We help you win.*

Architecture Alpha — the performance improvement that comes from better instructions to better-positioned advisors, working from a consolidated view of your entire capital structure.

About This Memo

This memo is provided for educational and informational purposes as a demonstration of CIOffice's thematic research perspective. It does not constitute investment advice, a recommendation, or an offer to buy or sell any security. CIOffice provides independent governance and capital architecture — all investment execution should be conducted through your regulated broker or bank.