THE AI ECOSYSTEM

Supply Chain Stack

A complete educational guide to the six layers of technology, infrastructure and energy that power artificial intelligence. The companies. The facts. The supply chain.

. . .

AN INFORMATION PACK BY MISS AI

For research and educational purposes only.
Not investment, financial, legal or tax advice.

INTRODUCTION

Why This Matters

Artificial intelligence does not run on one company. It runs on a complete ecosystem of six distinct layers of technology, infrastructure and energy. Each layer is as critical as the last. Remove any single layer and the entire AI supply chain stalls.

This guide breaks down every layer of that supply chain, the companies that operate inside it, what they do, how they make money, and why each layer matters. The aim is purely educational: to help you understand AI the way an industry analyst would, as a global supply chain.

IMPORTANT

This document is for general research and educational purposes only. It is not financial, investment, legal or tax advice and is not a recommendation to buy, sell or hold any company or product. See the full disclaimer at the end.

How This Pack Is Organised

LAYER 1

Chip Equipment Manufacturing

The machines that print the chips. Without them nothing else is possible.

LAYER 2

Chip Design

The architects who design the processors that train and run AI models.

LAYER 3

Chip Manufacturing, Foundries

The factories where chip designs are physically built into silicon.

LAYER 4

Memory

The high-speed memory that feeds AI chips fast enough to do their work.

LAYER 5

Data Centres

The physical real estate where AI chips are racked, cooled and connected.

LAYER 6

Power Grid & Energy

The electricity and infrastructure that keeps it all running 24/7.

LAYER 1 · CHIP EQUIPMENT MANUFACTURING

The Machines That Print the Chips

"You cannot print a chip without these machines."

Before a chip can be designed, before it can be manufactured, before it can power any AI system, someone has to build the machines that make chips possible. This is one of the most overlooked and most concentrated layers in the entire AI supply chain, and it is dominated by a single Dutch company.

LAYER 1 · COMPANY PROFILE

ASML Holding (ASML)

Listed	Nasdaq, Euronext Amsterdam
Headquarters	Veldhoven, Netherlands
Sector	Semiconductor Equipment
Ticker	ASML

What They Do

ASML makes lithography machines, the equipment that physically prints circuits onto silicon wafers to create computer chips. Their most advanced product is the EUV (Extreme Ultraviolet) lithography machine, the only tool on Earth capable of printing the smallest, most powerful chips that modern AI requires.

Without ASML machines, the chips inside every AI server, every smartphone and every advanced data centre cannot be made. There is no real alternative at the leading edge.

How the Machine Works

- The EUV machine fires a laser around 50,000 times per second at a tin droplet smaller than a human hair, generating ultraviolet light that etches circuits at a scale smaller than a virus.
- Each machine contains over 100,000 individual components, is the size of a double-decker bus and weighs more than 180 tonnes.
- Every unit must be disassembled, shipped in roughly 40 freight containers and reassembled on the customer's site.

The Concentrated Position

- ASML is currently the only company in the world that manufactures EUV lithography machines.
- Competitors including Nikon and Canon stepped back from leading-edge lithography over a decade ago.
- China has spent billions of dollars and more than ten years trying to replicate the technology and has not yet succeeded at the same scale.
- TSMC, Samsung and Intel have all publicly stated they cannot produce advanced chips without ASML.

LAYER 1 · ASML · THE NUMBERS

By the Numbers

EUR 32.7B
2025 total net sales (full year)

EUR 36-40B
2026 revenue guidance, raised in Q1 2026

EUR 38.8B
Order backlog at end of 2025

Up to EUR 380M
Approximate price of one EUV unit

60+
EUV Low-NA systems planned for 2026, 80+ for 2027

~52%
2025 gross margin, guided 51-53% for 2026

Geopolitical Importance

The US government pressured the Dutch government to ban ASML from selling its most advanced machines to China. ASML equipment now sits inside export-control regimes that treat it alongside military technology. The Dutch government coordinates with intelligence services to protect ASML facilities, and the company has become a flashpoint in US, European and Chinese tech policy.

What's Happening Right Now (2026)

- ASML raised its 2026 revenue guidance to EUR 36-40 billion in Q1 2026, citing strong AI infrastructure and memory chip demand.
- EUV revenue grew approximately 28% year-on-year in Q1 2026, including revenue from two High-NA systems.
- Hyperscaler AI capex is creating supply tightness; large memory orders from SK Hynix and Samsung are offsetting some logic-customer timing shifts.
- ASML plans to ramp Low-NA EUV output to at least 60 systems in 2026 and 80+ in 2027 to keep pace with demand.

LAYER 1 SUMMARY

Where it fits: every advanced AI chip starts here. Before any GPU is designed, before any fab is built, an ASML machine has to print the circuits.

LAYER 2 · CHIP DESIGN

The Architects of Artificial Intelligence

"Designs are worth more than the silicon they print on."

Chip designers create the blueprints for the processors that power AI. They do not build chips themselves. They design the architecture, the logic and the circuitry, then send those designs to foundries to be manufactured. This layer captures enormous value because the intellectual property in a chip design is worth far more than the physical silicon.

Four companies dominate this layer: NVIDIA, Broadcom, AMD and ARM. They compete in different ways. NVIDIA designs the dominant general-purpose AI GPU. Broadcom designs custom AI chips for individual hyperscalers. AMD is NVIDIA's primary direct competitor. ARM does not even design complete chips, it licenses the architecture that everyone else builds on top of.

LAYER 2 · COMPANY PROFILE

NVIDIA (NVDA)

Listed	Nasdaq
Headquarters	Santa Clara, California, USA
Sector	Semiconductors / AI Computing
Ticker	NVDA

What They Do

NVIDIA designs GPUs (Graphics Processing Units) that have become the dominant hardware for training and running AI models. Originally built for gaming graphics, NVIDIA's chips turned out to be well suited for the parallel processing AI demands. Today, virtually every major AI model, including ChatGPT, Gemini, Grok and Claude, was trained on NVIDIA hardware.

Market Position

- NVIDIA holds approximately 90% market share in AI data-centre accelerators.
- Its H100, H200, Blackwell and Blackwell Ultra GPUs are the standard hardware for AI infrastructure globally.
- Microsoft, Google, Amazon, Meta, Oracle and every major AI lab are NVIDIA customers.
- NVIDIA's CUDA software platform is the AI developer standard. Switching away from CUDA would mean rewriting an enormous installed base of code.

By the Numbers (Fiscal 2026)

$215.9B
Full-year FY2026 revenue, up 65% year-on-year

$193.7B
FY2026 data centre revenue

$68.1B
Q4 FY2026 quarterly revenue, up 73% year-on-year

$62.3B
Q4 FY2026 data centre revenue, up 75% year-on-year

75%
GAAP gross margin in Q4 FY2026

~$500B
Visibility on Blackwell + Rubin revenue through end of 2026

LAYER 2 · NVIDIA · WHAT'S NEW

The Story Right Now

Blackwell Is the Engine

Blackwell and Blackwell Ultra are now in volume production and deployed across all major cloud providers. Demand continues to outstrip supply, and Blackwell GB300 shipments have crossed GB200, now accounting for roughly two-thirds of Blackwell revenue. Networking inside the data centre, the chips and switches that connect GPUs together, reached $11 billion in Q4 FY2026 alone, more than 3.5x the prior year.

Rubin Is Already Lined Up

NVIDIA has unveiled the Rubin platform, comprising six new chips that aim to deliver up to a 10x reduction in inference token cost compared with Blackwell. AWS, Google Cloud, Microsoft Azure and Oracle Cloud Infrastructure are among the first announced deployers of Vera Rubin instances.

Mega-Deals and Sovereign AI

- NVIDIA announced a strategic partnership with OpenAI to deploy at least 10 gigawatts of NVIDIA systems for OpenAI's next-generation AI infrastructure.
- Anthropic announced it will run and scale on NVIDIA infrastructure, initially adopting around 1 gigawatt of compute capacity using Grace Blackwell and Vera Rubin systems.
- The first Blackwell wafer was produced on US soil at TSMC's Arizona facility.
- The US Department of Energy is building Solstice, its largest AI supercomputer, with Oracle and 100,000 NVIDIA Blackwell GPUs.

China Headwinds

Export restrictions continue to weigh on the China business. NVIDIA has been effectively shut out of the mainland AI compute market for now, and is guiding to no data-centre compute revenue from China in upcoming quarters, an unusually explicit acknowledgement of geopolitical risk.

LAYER 2 · COMPANY PROFILE

Broadcom (AVGO)

Listed	Nasdaq
Headquarters	Palo Alto / San Jose, California, USA
Sector	Semiconductors / Networking / Software
Ticker	AVGO

What They Do

Broadcom designs custom AI chips (ASICs) and networking semiconductors. While NVIDIA sells general-purpose GPUs, Broadcom builds bespoke chips for individual customers, chips optimised for exactly what Google, Meta or other hyperscalers need their AI models to do. Their networking chips are also the backbone that connects servers inside virtually every major data centre on Earth.

The Custom Chip Story

- Broadcom builds Google's TPUs (Tensor Processing Units), custom AI chips for Meta, and networking silicon for almost every hyperscaler.
- Custom chips can be more efficient than general-purpose GPUs for specific AI tasks, and at lower cost per workload.
- Broadcom is a major reason hyperscalers are not 100% dependent on NVIDIA.
- Without Broadcom networking, individual chips inside a data centre cannot communicate fast enough to train very large AI models.

Why It Matters in the Supply Chain

Broadcom sits at two megatrend points at once. As hyperscalers spend hundreds of billions of dollars on AI capex, an increasing share is going to custom silicon, which favours Broadcom, and the rest still needs Broadcom networking gear to tie the boxes together. AI ASIC revenue is the fastest-growing segment of the company.

Talking Points (2026)

- Broadcom's AI custom-silicon business is reportedly in advanced discussions with multiple new hyperscale customers.
- Networking silicon attached to AI clusters is now one of the largest growth lines in the entire semiconductor industry.
- Broadcom continues to pay a growing dividend.

LAYER 2 · COMPANY PROFILE

AMD (AMD)

Listed	Nasdaq
Headquarters	Santa Clara, California, USA
Sector	Semiconductors
Ticker	AMD

What They Do

AMD designs CPUs and GPUs and is NVIDIA's primary like-for-like competitor in AI accelerators. Its Instinct MI300X and MI350 series chips are winning meaningful contracts at Microsoft, Meta and Oracle as customers seek alternatives to a single-source supply. AMD also dominates roughly half of the global data-centre CPU market alongside Intel.

The Competitive Position

- MI300X and MI350 are gaining traction as alternatives to NVIDIA, especially for AI inference workloads.
- Microsoft and Meta have publicly deployed AMD chips at scale, reducing their NVIDIA dependency.
- AMD's MI400 / MI430X / MI440X / MI455X roadmap and Helios rack-scale architecture target the next generation of frontier AI training.
- Hyperscalers prize having a credible second-source supplier of AI accelerators.

CEO & Strategy

AMD's CEO Lisa Su is widely regarded as one of the strongest executives in semiconductors. She turned AMD around from near-bankruptcy in 2014 into a credible challenger across CPUs, GPUs and AI accelerators, a track record that has earned the company significant credibility with hyperscale buyers who want a second source.

Talking Points (2026)

- AMD's AI accelerator revenue continues to grow at triple-digit rates from a smaller base.
- Hyperscaler diversification is the structural force behind AMD's AI accelerator demand.
- Server CPU share gains versus Intel continue, providing a more stable cash-flow base under the AI business.

LAYER 2 · COMPANY PROFILE

ARM Holdings (ARM)

Listed	Nasdaq
Headquarters	Cambridge, United Kingdom
Sector	Semiconductor IP / Architecture
Ticker	ARM

What They Do (and Don't Do)

ARM does not make chips. ARM does not even design complete chips end-to-end. ARM designs the architecture, the fundamental instruction set and core designs that other companies license to build their own chips. It is one of the most distinctive business models in technology.

The Ubiquity

- ARM-based chip designs are inside roughly 99% of all smartphones on Earth.
- Hundreds of billions of ARM-based chips have been shipped since the company was founded.
- Every iPhone, every Android phone, every Apple Mac and every Amazon Alexa runs on ARM architecture.
- NVIDIA tried to acquire ARM for $40 billion. Regulators blocked it, which is itself a measure of how strategically critical ARM is.

The Royalty Business Model

ARM charges a licensing fee when a company uses its architecture to design a chip, then takes a royalty on every chip sold using that design. That means ARM earns money every time Apple sells an iPhone, every time a server ships, every time any ARM-based device is manufactured, without building anything physical itself. It is one of the purest royalty businesses in technology.

The AI Expansion

- Microsoft, Google and Amazon are all building custom ARM-based server CPUs for their data centres, each one a new royalty stream.
- Apple's M-series chips, which now power Macs and iPads, are ARM-based and increasingly run on-device AI.
- ARM is extending its Neoverse data-centre architecture specifically for AI inference and AI-class workloads.
- Royalty revenue from AI-related chips is one of the company's fastest-growing lines.

LAYER 3 · CHIP MANUFACTURING (FOUNDRIES)

The Factories the World Cannot Live Without

"Designs are nothing without a foundry that can build them."

Chip designers create blueprints. Foundries build the actual physical chips. This requires some of the most complex, expensive and precise manufacturing processes in human history. The companies in this layer own and operate the fabs, fabrication facilities, where silicon wafers are transformed into finished chips. One company dominates the leading edge.

LAYER 3 · COMPANY PROFILE

TSMC, Taiwan Semiconductor Manufacturing Co. (TSM)

Listed	NYSE (ADR), Taiwan Stock Exchange
Headquarters	Hsinchu, Taiwan
Sector	Semiconductor Manufacturing (Foundry)
Ticker	TSM

What They Do

TSMC is the world's most important manufacturer of advanced semiconductor chips. They do not design chips. They build chips designed by others. Apple brings them iPhone designs. NVIDIA brings them GPU designs. AMD brings them CPU and accelerator designs. TSMC builds all of them, better than anyone else on Earth at the leading edge.

What a Fab Actually Is

A fab (fabrication facility) is the factory where chips are physically made. TSMC's fabs are unlike any other factory in existence:

- The air inside is roughly 1,000 times cleaner than a hospital operating theatre. A single dust particle can destroy a chip.
- Workers wear full bunny suits. No skin, hair or breath can be exposed to the manufacturing environment.
- The machines inside cost hundreds of millions of dollars each.
- Building a single advanced fab costs $10-20 billion and takes 3-5 years.
- TSMC operates fabs the size of small cities, some covering more than 200,000 square metres.

LAYER 3 · TSMC · THE NUMBERS

Scale, Spend, and the Silicon Shield

~$122B
2025 net revenue, up roughly 36% year-on-year

$33.7B
Q4 2025 revenue, up 25% year-on-year

>90%
Estimated share of the world's most advanced chips

~$165B
Total committed Arizona investment program

Up to 11
Arizona fab phases announced in 2026 framework

Sold out
Advanced node capacity through 2026 and into 2027

The Customer List

Every NVIDIA AI GPU is manufactured by TSMC. Apple's entire chip supply, from iPhones to MacBooks, is made by TSMC. AMD's MI-series accelerators, Broadcom's custom ASICs, and many of Google's, Amazon's and Meta's in-house silicon designs are all manufactured by TSMC. Apple, AMD, Qualcomm, Broadcom and NVIDIA all manufacture chips at TSMC's Arizona site.

The Silicon Shield

Taiwan is claimed by China, and TSMC sits at the centre of the resulting geopolitical tension. Taiwan calls TSMC its "Silicon Shield": the theory being that global dependence on TSMC makes any conflict prohibitively costly for everyone involved. If Taiwan-based capacity were disrupted, the global tech economy would effectively halt within months.

That is why every major economy is paying TSMC to build inside its borders. The US allocated approximately $6.6 billion in CHIPS Act grants plus billions in loans to support TSMC's Arizona complex. Japan and Germany have backed similar projects. Under a 2026 US-Taiwan tariff framework, TSMC's Arizona program has been broadened to up to 11 fab phases over a multi-decade rollout.

What's Happening Right Now (2026)

- Arizona Fab 1 is producing 4nm chips for Apple, NVIDIA, AMD, Qualcomm and Broadcom. The first Blackwell wafer was produced there on US soil.
- Fab 2 is complete with tool installation in 2H 2026 and 3nm volume production targeted for 2027. Advanced node capacity is fully booked through 2026.

LAYER 4 · MEMORY

The Fuel That AI Runs On

"A GPU without HBM is a Ferrari without fuel."

AI chips are extraordinarily powerful processors, but they need memory to feed them data fast enough to be useful. Standard computer memory is too slow. AI requires a specialised, high-performance memory called HBM (High Bandwidth Memory), and only three companies in the world produce it at scale.

HBM stacks memory chips vertically, like a skyscraper, and connects them with thousands of tiny wires, creating a superhighway for data. This lets AI accelerators access memory many times faster than traditional approaches and is the reason modern AI training is even possible.

MARKET CONTEXT

Demand for HBM is so intense that DRAM prices have risen sharply in 2026 because HBM is "crowding out" conventional memory capacity. Micron has noted that HBM consumes roughly three wafers for every one wafer of equivalent DDR5 capacity.

LAYER 4 · COMPANY PROFILE

SK Hynix (000660.KS)

Listed	Korea Stock Exchange (KOSPI)
Headquarters	Icheon, South Korea
Sector	Semiconductor Memory
Ticker	000660.KS

What They Do

SK Hynix is the world's leading producer of HBM, the specialised memory chips that sit directly alongside NVIDIA's AI GPUs and allow them to process data fast enough to run modern AI models. SK Hynix has been the primary memory supplier for every major NVIDIA AI generation: A100, H100, H200 and the Blackwell family.

The NVIDIA Connection

- SK Hynix held an estimated 62% share of the HBM market in Q2 2025 and around 57% by revenue in Q3 2025.
- It has met NVIDIA's certification at every recent generation of HBM.
- NVIDIA's CEO Jensen Huang has personally engaged with SK Hynix leadership to secure supply, a measure of how critical this relationship is.
- UBS forecasts SK Hynix could reach approximately 70% share of HBM4 sold into NVIDIA's next-generation Rubin platform.

Where Things Stand in 2026

- HBM3E is expected to remain the dominant HBM product through 2026, accounting for roughly two-thirds of HBM shipments while HBM4 ramps.
- SK Hynix has completed development of HBM4 and is the world's first supplier with a mass-production system in place.
- It is also expected to be the first HBM3E supplier for Google's seventh-generation TPU (v7p / v7e).
- Capacity is fully committed; SK Hynix is building the Cheongju M15X fab specifically to expand HBM3E and HBM4 output.
- Samsung and SK Hynix have raised HBM3E supply prices for 2026 by roughly 20% on the back of strong demand.

LAYER 4 · COMPANY PROFILE

Micron Technology (MU)

Listed	Nasdaq
Headquarters	Boise, Idaho, USA
Sector	Semiconductor Memory
Ticker	MU

What They Do

Micron is the only US-headquartered manufacturer of DRAM, NAND flash and HBM at scale. As the AI memory market expands, Micron sits in a strategically important position: it has US government support via the CHIPS Act, growing HBM market share, and a domestic-champion role that makes it politically important well beyond its size.

The Story

- Micron overtook Samsung to become the #2 HBM supplier, holding around 21% market share in Q2 2025.
- It is shipping HBM3E to NVIDIA and AMD's MI350 family, and has begun shipping 11 Gbps HBM4 samples.
- Micron has announced significant US capacity expansion plans backed by federal CHIPS Act funding.
- It claims roughly 30% lower power consumption versus competing HBM3E offerings, a meaningful efficiency edge.

Why It Matters in the Supply Chain

Memory has historically been a cyclical, low-margin commodity. The AI cycle is changing that. With HBM crowding out conventional DRAM and demand from hyperscalers, custom-ASIC programs (Google TPU, Amazon Trainium) and NVIDIA all hitting at once, the supply side is structurally tight. Micron is the largest US-based memory supplier exposed to that dynamic.

Talking Points (2026)

- Memory pricing is in an unusually strong cycle. DRAM and HBM3E prices have risen sharply through 2025-2026.
- Micron is competing with Samsung and SK Hynix for NVIDIA HBM4 and the upcoming 16-Hi HBM4 supply requested for late 2026.
- US government support via the CHIPS Act is a strategic factor that the Korean competitors do not enjoy on the same terms.

LAYER 4 · COMPANY PROFILE

Samsung Electronics (005930.KS)

Listed	Korea Stock Exchange (KOSPI)
Headquarters	Suwon, South Korea
Sector	Semiconductors / Consumer Electronics / Memory / Foundry
Ticker	005930.KS

What They Do

Samsung is simultaneously the world's largest consumer electronics company and one of the three dominant memory chip manufacturers. It produces DRAM, NAND and HBM, while also operating its own foundry business that competes (at a distance) with TSMC for advanced chip manufacturing.

Memory Position

- Samsung holds roughly 40% of the global DRAM market across all categories.
- In HBM specifically it slipped to third place behind SK Hynix and Micron, with around 17% share in Q2 2025.
- Samsung has now passed NVIDIA's qualification for 12-layer HBM3E, a symbolic recovery of credibility.
- Counterpoint Research expects Samsung's HBM share to recover above 30% as HBM4 ramps in 2026.

The HBM4 Battleground

The competitive battle is shifting to HBM4, the next-generation memory designed to launch with NVIDIA's Rubin platform. Samsung has demonstrated up to 11 Gbps signalling, ahead of SK Hynix's 10 Gbps reference, and is in HBM4 talks with NVIDIA, Broadcom and Google. Volume HBM4 supply is expected to begin in the first half of 2026.

Other Considerations

- Samsung's foundry continues to compete with TSMC for advanced chip orders, including for its own custom ASICs and external customers.
- Diversified consumer-electronics revenue gives Samsung a more cyclical profile than the pure-play memory names but a much larger absolute earnings base.
- It is one of the few single companies with exposure to memory, foundry and consumer AI hardware.

LAYER 5 · DATA CENTRES

The Real Estate of the AI Economy

"AI chips need somewhere to live, and the world is running out of suitable buildings."

Training an AI model requires thousands of chips running simultaneously in climate-controlled, ultra-secure, power-hungry facilities. Data centres are the physical infrastructure of the AI economy, and demand is growing faster than the global construction industry can deliver supply.

LAYER 5 · COMPANY PROFILE

Equinix (EQIX)

Listed	Nasdaq
Headquarters	Redwood City, California, USA
Sector	Data Centre REIT
Ticker	EQIX

What They Do

Equinix is the world's largest data centre company, structured as a Real Estate Investment Trust (REIT). It owns and operates the physical buildings where the internet, cloud computing and AI infrastructure actually live. Microsoft, Google, Amazon, Meta and Apple all pay Equinix to house servers and to interconnect with each other inside its facilities. In effect, Equinix is the landlord of the internet.

Scale

- 260+ data centres across more than 70 cities in over 30 countries.
- Tens of thousands of customers globally, including over a thousand cloud and network providers interconnecting inside its buildings.
- Equinix sites sit at the most critical internet exchange points on Earth, where global networks physically connect.

Business Model

- REIT structure. Legally required to distribute most taxable income as dividends to shareholders.
- Customer contracts typically run 3-10 years, making revenue highly predictable and 'sticky'.
- Once a customer racks servers in an Equinix facility, the cost and risk of moving out is so high that they almost never do.
- Long track record of consecutive revenue growth, one of the longest growth streaks of any US-listed company.

Why AI Changes the Equation

AI workloads consume far more power per server rack than traditional cloud workloads, often 10x or more. That requires entirely different cooling, power and design specifications. Equinix is one of the few operators with the global footprint, customer base and capital to retrofit for liquid cooling and to build new AI-specific facilities. AI demand is, in effect, creating a data centre capacity crunch.

LAYER 5 · COMPANY PROFILE

Microsoft (MSFT), AI Infrastructure

Listed	Nasdaq
Headquarters	Redmond, Washington, USA
Sector	Software / Cloud Infrastructure
Ticker	MSFT

What They Do (in this layer)

Microsoft is best known for software, but it is also one of the largest builders and operators of AI data centres in the world. Its Azure cloud is the second-largest cloud platform globally and is the primary commercial deployment platform for OpenAI's products, including ChatGPT. Microsoft's data centre and energy commitments are arguably more important to its AI economics than any single chip purchase.

Scale of the Build

- Microsoft committed roughly $80 billion in 2025 to build new AI data centres, more than the annual GDP of many countries.
- Azure is the primary commercial channel for OpenAI, which means OpenAI demand effectively flows into Microsoft's infrastructure capex plans.
- Microsoft has signed long-term, fixed-price power purchase agreements specifically to lock in 24/7 clean electricity for AI workloads.

Why Microsoft Belongs in Both the Data Centre and Energy Discussion

Microsoft has effectively become a quasi-utility customer. Its 20-year deal to bring back the Three Mile Island Unit 1 reactor (rebranded the Crane Clean Energy Center) is the first time a retired US nuclear plant has been brought back to life specifically for one commercial customer. It is a window into how the largest AI companies are now thinking about electricity supply as a structural input cost, not just an operating expense.

WHY IT MATTERS

Microsoft's infrastructure commitments are large enough to influence the data centre, electricity and even the nuclear industries simultaneously.

LAYER 6 · POWER GRID & ENERGY

The Hidden Bottleneck

"AI does not have a chip problem any more. It has an electricity problem."

AI has a quiet secret: it is extraordinarily power-hungry. Training a single large AI model can consume as much electricity as hundreds of homes use in a year. Running AI models at scale 24/7 across thousands of data centres requires a level of generation and grid infrastructure that the world is simply not yet ready for. This is the most underbuilt and most urgent constraint in the entire AI supply chain.

The Electricity Crunch in Numbers

~8%
Estimated share of US electricity that AI data centres could consume by 2030

~$2.5T
Estimated US grid investment needed over the next decade

18-24 mo
Current waiting list for power transformers

Multi-GW
Scale of single-customer power deals being signed in 2026

Major hyperscalers are reportedly delaying data centre openings because they cannot get enough electricity connected to the grid fast enough. The constraint has shifted from chips to kilowatts.

LAYER 6 · COMPANY PROFILE

Constellation Energy (CEG)

Listed	Nasdaq
Headquarters	Baltimore, Maryland, USA
Sector	Nuclear / Utilities
Ticker	CEG

What They Do

Constellation Energy is America's largest nuclear energy producer, operating a fleet of around 21 reactors with approximately 22 gigawatts of total nuclear capacity. Nuclear power is emerging as the preferred energy source for AI data centres because it provides 24/7, carbon-free, uninterruptible electricity, exactly what AI workloads demand.

The Microsoft Deal

- Constellation signed a 20-year power purchase agreement with Microsoft to restart the Three Mile Island Unit 1 reactor, rebranded the Crane Clean Energy Center.
- It will add approximately 835 MW of carbon-free electricity to the PJM grid, enough to power around 800,000 homes.
- It is the first time a fully retired US nuclear plant has been brought back into service specifically for a single commercial customer.
- The plant is currently planned to come online in the 2027-2028 window, with regulatory approvals and refurbishment underway.
- Constellation expects the deal to lift its base earnings-per-share growth target from at least 10% to at least 13% per year through 2030.

The Bigger Strategic Picture

Constellation is also pursuing the $26.6 billion acquisition of Calpine and signing additional long-term clean power deals with data-centre operators. Together with the Three Mile Island restart, this positions the company as the largest single source of contracted, carbon-free baseload power for the AI buildout in the United States.

WHY IT MATTERS

If chips are the brains of AI, electricity is the bloodstream. Constellation is one of the largest contracted suppliers of the right kind of electricity for the job.

FRAMEWORK

How To Use This Pack

Think In Layers, Not Tickers

The single most useful idea in this pack is that AI is not a stock, it is a supply chain. Each layer has different economics, different risks and different sensitivities to the cycle. Understanding the layers separately gives you a much clearer picture of where the value is being created and where the bottlenecks are than looking at any single name in isolation.

Map Each Layer To A Question

Layer 1, Equipment

Who is currently the only company that can sell the leading-edge chip-printing tools? ASML.

Layer 2, Design

Who is paid for every advance in compute, and who is paid for every chip ever shipped? NVIDIA, AMD, Broadcom and ARM.

Layer 3, Foundry

Who actually manufactures the silicon at the leading edge? TSMC.

Layer 4, Memory

Who is solving the memory bottleneck that constrains every AI chip? SK Hynix, Micron and Samsung.

Layer 5, Data Centres

Who owns or builds the physical infrastructure where the chips actually live? Equinix and the hyperscalers, led by Microsoft.

Layer 6, Energy

Who supplies the 24/7 carbon-free electricity that the AI buildout depends on? Constellation Energy and the broader nuclear and grid complex.

Where The Bottlenecks Sit

In any supply chain, the most strategically important points tend to be where capacity is tightest. In the AI supply chain today, the obvious tight points are EUV lithography, leading-edge foundry capacity, HBM memory, AI-grade data centre real estate, and 24/7 carbon-free electricity. Every one of those is currently sold out or capacity-constrained, which is what makes them so important to understand.

And The Risks To Be Aware Of

- Geopolitical risk, especially Taiwan, China export controls and US-EU tariff dynamics.
- Cyclical risk, memory in particular has historically been highly cyclical.
- Regulatory risk, antitrust, energy permitting and nuclear regulation can all move slowly.
- Concentration risk, a small number of customers (the hyperscalers) drive most of this demand.
- Technology risk, new architectures (custom silicon, optical compute, quantum) can shift the balance between layers.

REMEMBER

This is purely a description of how the AI supply chain is structured today. Nothing here is a recommendation, a strategy, or advice of any kind. Please read the full disclaimer on the following pages.

FINAL SECTION

Important Disclaimer

Please read this disclaimer carefully before relying on any information in this pack.

1. Educational and Informational Purpose Only

This document, and all information contained within it, is provided strictly for general research, educational and informational purposes. It is not intended to be, and must not be construed as, financial advice, investment advice, trading advice, legal advice, tax advice, accounting advice, retirement planning advice, or any other form of professional advice. Nothing in this document constitutes a recommendation, suggestion, endorsement or solicitation to buy, sell, hold, subscribe for, short, or otherwise transact in any security, financial instrument, asset, derivative, cryptocurrency, fund, or product of any kind, in any jurisdiction. References to companies, sectors, products and supply chains are descriptive and educational only.

2. No Personalised Advice and No Fiduciary Relationship

The content of this pack has been prepared without reference to the specific financial situation, objectives, tax position, risk tolerance, time horizon, regulatory status, or particular needs of any individual reader. By providing this pack, the author does not enter into any advisory, fiduciary, agency, or client relationship with you. You must not rely on the content as a substitute for advice from a qualified, licensed financial adviser, tax adviser, lawyer, accountant, or other professional in your own jurisdiction. Before making any financial, business or other decision relating to anything described here, you should obtain independent professional advice that takes your individual circumstances into account, and where relevant read the official offer documents, prospectuses or disclosure documents associated with any product you may be considering.

3. Not An Offer or Solicitation

Nothing in this pack constitutes or forms part of any offer, solicitation, invitation, inducement, promotion, marketing communication, or recommendation to buy, sell, subscribe for, or otherwise deal in any security or financial product, in any jurisdiction in which such an offer, solicitation, invitation, inducement, promotion, marketing communication, or recommendation would be unlawful. The companies, products and supply chains mentioned in this pack are referenced solely as examples of how the global AI supply chain is structured. No commercial, endorsement, sponsorship, partnership, affiliate or other relationship between the author and any company referenced should be inferred unless expressly stated.

4. Forward-Looking Statements and Past Performance

This pack contains statements about future events, projections, expectations, opinions, market outlooks, and forward-looking commentary, including statements about industry growth, capital expenditure plans, technology roadmaps, and the strategies of third-party companies. Forward-looking statements are inherently uncertain and depend on numerous assumptions, market conditions, regulatory developments, technological developments, geopolitical events, and other factors that are outside any one party's control. Actual outcomes may differ materially from any expectations expressed or implied. Past performance, historical data and prior trends are not reliable indicators of future results.

FINAL SECTION · CONTINUED

Important Disclaimer (continued)

5. No Warranty as to Accuracy or Completeness

While reasonable care has been taken in preparing this pack, the information has been compiled from a variety of public sources, including company filings, news media and industry research, and the AI sector evolves rapidly. The author makes no representation, warranty, guarantee or undertaking, whether express or implied, as to the accuracy, completeness, currency, reliability, fitness for purpose or non-infringement of any information in this pack. Numbers, valuations, market shares, company strategies, regulatory approvals and product roadmaps may change after the date of preparation and may already be out of date by the time you read this. You are solely responsible for verifying any information that you intend to rely on.

6. Risk Warning

Any decision relating to securities of any company referenced in this pack involves significant risk, including the risk of partial or total loss of capital. Risks include, but are not limited to: market risk, liquidity risk, currency and exchange-rate risk, geopolitical risk, regulatory and tax risk, technology risk, supply-chain risk, concentration risk, counterparty risk, and the specific business risks of each company. The semiconductor, AI infrastructure, data centre and energy sectors discussed here are particularly cyclical, capital-intensive, and exposed to rapid technology change. Any reader considering a financial decision in connection with any of these topics should consult a qualified professional in their own jurisdiction first.

7. Jurisdiction and Regulatory Notice

This pack is intended for a global audience and is not directed at any particular country, jurisdiction or regulatory regime. The author is not a registered or licensed financial adviser, broker-dealer, investment adviser, fund manager, or financial services provider in any jurisdiction, and does not provide regulated financial services. The information in this pack may not be appropriate, lawful, or available for use in all jurisdictions. It is the reader's responsibility to ensure that any use of this information complies with all applicable laws, regulations and rules in their own country or jurisdiction. If access to this content would be unlawful in your jurisdiction, you must not access or rely on it.

8. No Liability

To the maximum extent permitted by applicable law, the author and any associated entities, officers, employees, contractors and agents disclaim all liability for any loss, damage, cost or expense (whether direct, indirect, incidental, consequential, special, exemplary, punitive or otherwise) arising from or in connection with: (a) any reliance placed on the information in this pack; (b) any errors, omissions or inaccuracies in the content; (c) any actions taken or not taken as a result of reading this pack; or (d) any decisions made in respect of any security, asset, business strategy or other matter referenced. Your use of this pack is entirely at your own risk.

9. Third-Party Names and Trademarks

All company names, product names, brand names, ticker symbols and trademarks mentioned in this pack are the property of their respective owners. They are referenced for identification, research, commentary and educational purposes only. No endorsement, sponsorship, certification, approval, partnership or affiliation by or with any of the companies named is claimed or implied.

10. Acknowledgement

By reading, downloading, sharing or otherwise using this pack, you acknowledge that you have read, understood and accepted this disclaimer in full. If you do not agree with any part of this disclaimer, you must not rely on the content of this pack and should discontinue use of it immediately.