The Foundation of AI

How NVIDIA, ASML and TSMC underwrite the AI economy

Keeley Walsh | Associate Analyst, Research

26 May 2026

The foundation of the AI infrastructure stack

Three companies. One supply chain. The entire AI build-out.

As we prefaced in our most recent AI education series, "The AI Economy", almost every AI breakthrough you read about ChatGPT, Gemini, Claude, autonomous driving stacks and drug discovery models, runs on chips designed by NVIDIA, manufactured by TSMC, and printed by ASML's lithography machines.

Why this matters

If you strip any one of them, the AI economy stalls. There is no second source for these critical infrastructure inputs. It is important to understand this structural backdrop relating to AI exposure, semiconductor cyclicality, and concentration risk in the global tech complex.

~80%

of data-centre spend in 2026 runs through NVIDIA chips.

>90%

of advanced (sub-7nm) wafers in the world are fabricated by TSMC.

100%

of EUV lithography systems on Earth are built by ASML. No competitor exists.

Source: Company data, Morgans estimates

Three layers of the infrastructure stack

The three companies at an individual level

NVIDIA
The designer

HEADQUARTERS
Santa Clara, USA

JOB: DESIGN
GPUs, accelerators, networking, CUDA software

FLAGSHIP PRODUCT
Blackwell GPU (TSMC 4NP), Rubin GPU (TSMC N3, late 2026)

AI ROLE
The brain of every training and inference cluster

TSMC
The manufacturer

HEADQUARTERS
Hsinchu, Taiwan

JOB: MANUFACTURE
Contract manufacturing fabricating the chips others (Nvidia, Apple, AMD, Broadcom, Qualcomm) design

FLAGSHIP PRODUCT
N3, N2 (2026), A16 and CoWoS advanced packaging manufacturing processes

AI ROLE
Only fab on Earth that can build leading-edge AI silicon at scale

ASML
The equipment

HEADQUARTERS
Veldhoven, Netherlands

JOB: PROVIDE EQUIPMENT
Photolithography machines that print chip features

FLAGSHIP PRODUCT
Low-NA EUV ($400m+), High-NA EUV ($350m+) lithography machines

AI ROLE
Sole supplier of the equipment that makes sub-7nm possible

Source: Company data, Morgans estimates

NVIDIA: The Designer

Layer 01

What NVIDIA actually does

- NVIDIA designs the GPUs and accelerators that train and run almost all serious AI models. It does not manufacture them.
- Its real moat is software. The CUDA programming framework, libraries, and developer ecosystem mean that nearly all AI research is written against NVIDIA hardware by default.
- Product roadmap: Blackwell (current generation, TSMC 4NP) → Rubin (Q3 2026, TSMC N3, ~336 billion transistors) → Rubin Ultra and Vera platforms beyond. These are like iPhone models for chips.

Why it is non-negotiable in the AI stack

- Roughly 80% of data-centre AI accelerator revenue runs through NVIDIA in 2026.
- The hyperscalers - Microsoft, Amazon, Google, Meta, Oracle - build their AI fleets around its chips.
- Switching costs are extreme. Porting models, kernels and tooling off CUDA is theoretically possible but operationally painful and expensive to execute. Think switching from an android to an iPhone.

Key company highlights

FY26/27F data-centre revenue from selling chips

~US$194bn/US$335bn

Market share of AI chips

~80-85%

Nvidia's software tools that provide a moat

CUDA, cuDNN, TensorRT

Manufacturing

100% outsourced to TSMC

Advanced packaging

CoWoS also done by TSMC

Hyperscaler dependency

Microsoft, Amazon, Google, Meta

Source: Company data, Visible Alpha, Morgans estimates

TSMC: The Manufacturer

Layer 02

What TSMC actually does

- TSMC is a pure-play contract foundry. Customers (NVIDIA, Apple, AMD, Broadcom, Qualcomm) bring designs and TSMC builds them.
- It dominates the leading edge. ~74% of TSMC's wafer revenue now comes from 7nm and below, with 3nm alone at 25%.
- It also owns CoWoS (Chip-on-Wafer-on-Substrate) the advanced packaging that bonds AI logic to high-bandwidth memory. There is no comparable alternative at volume. Right now demand so far exceeds supply that it is the single biggest reason AI chips are delayed or short in supply.

Why it is non-negotiable in the AI stack

- Taiwan accounts for over 90% of the world's most advanced chip production. The runners-up being Samsung Foundry and Intel Foundry are one to two nodes behind on yields.
- NVIDIA has reportedly reserved over 60% of TSMC's entire 2026 CoWoS output. The constraint, not the design, is what gates AI shipment volumes.

Key company highlights

Global sub-7nm share (most advanced powerful chips)

>90% globally

Current and upcoming chip generations

N3, N2 (2026), A16 (2027)

Key bottleneck

CoWoS advanced packaging

CoWoS ramp to meet current demand

35k → 120–140k wafers/mo by end-2026

Largest customers

Apple, NVIDIA, AMD, Broadcom

Equipment

ASML for EUV, no substitute

Source: Company data, Morgans estimates

ASML: The Equipment

Layer 03

What ASML actually does

- ASML builds photolithography machines which is the equipment that uses light to print transistor patterns onto silicon wafers.
- Its EUV (extreme ultraviolet) machines are the only tools on Earth capable of printing features small enough for 7nm and below. Each system costs US$200m+ for Low-NA, US$380m+ for High-NA.
- An EUV machine is arguably the most complex industrial product ever built. It is assembled from 100,000+ components from 800+ suppliers and shipped on 40 freight containers.

Why it is non-negotiable in the AI stack

- ASML has a 100% global monopoly on EUV. Nikon and Canon abandoned EUV development a decade ago.
- Without EUV, sub-7nm chips simply cannot be made at commercial yields. Without sub-7nm, there is no Blackwell (Nvidia's newest, no Rubin, no current-generation AI training cluster).
- Delivery lead times now stretch 12-18 months.

Key company highlights

EUV market share
100% (monopoly)

Low-NA EUV (current standard machine)
60 systems planned in 2026

High-NA EUV (newer, more powerful machine)
Intel, Samsung first customers; TSMC ~2029

Critical supplier
Germany company Zeiss SMT is the sole EUV optics maker

2026 revenue guidance
€34–39bn

Largest customer
TSMC

Source: Company data, Morgans estimates

How one chip is made and who needs whom

The interdependency between the three companies

1 ASML

ships EUV lithography machine to a TSMC fab

2 TSMC

uses ASML's EUV to fabricate the NVIDIA chip

3 NVIDIA

ships GPU systems to hyperscalers and labs

The lithography lock

- Chips are made by "printing" tiny circuits - ASML makes the only machines capable of printing circuits small enough for modern AI chips. No other company has figured out how to build them, and it took ASML 20+ years to do it.
- This makes ASML an irreplaceable bottleneck - no ASML machines means no advanced chips, full stop. Each machine costs up to $380 million and takes over a year to deliver.
- TSMC and ASML are completely locked together - as TSMC pushes to make chips smaller and faster, they can only move as fast as ASML can deliver the next generation of machines.

The manufacturing lock

- NVIDIA only designs chips - it doesn't build them. All of NVIDIA's AI chips are manufactured exclusively by TSMC, and NVIDIA has locked up the majority of TSMC's most critical production capacity for 2026. No other manufacturer (Samsung, Intel) is advanced enough to substitute.
- When AI chips are delayed or short in supply, TSMC is almost always the bottleneck - not NVIDIA. The constraint is physical manufacturing capacity, not NVIDIA's ability to design.

Source: Company data, Morgans estimates

The circular web

The relationships are not a one-way chain. Each company underwrites the others' future demand

Does ASML depend on NVIDIA?

Yes, indirectly but heavily.

- AI demand is what makes TSMC, Samsung, and Intel willing to spend US$300m+ per EUV machine.
- Roughly half of ASML's annual orders are tied directly to AI-driven capex.
- If AI capex cools, ASML's order book is the first thing to compress.

Does TSMC depend on NVIDIA?

Increasingly, and dangerously so.

- NVIDIA is now TSMC's second-largest customer behind Apple, and is closing fast.
- Over 60% of 2026 CoWoS capacity is reserved for NVIDIA.
- Customer concentration risk means a single down-cycle at NVIDIA would leave TSMC packaging lines idle.

Does NVIDIA depend on ASML?

Yes, one step removed.

- NVIDIA buys nothing from ASML directly, but every NVIDIA chip is printed on ASML equipment.
- Delays in ASML's High-NA roadmap eventually constrain NVIDIA's ability to scale beyond Rubin Ultra.
- ASML is, effectively, the upstream gatekeeper of NVIDIA's long-term performance curve.

Source: Company data, Morgans estimates

What if one of them stumbles?

Triggers and consequences

ASML

TRIGGERS

- Zeiss (German input supplier) optics defect
- Plant disruption
- Export controls
- High-NA (new EUV machine) delays

CONSEQUENCE

Existing fabrications will keep producing, there is no instant cliff, however within 12-18 months, the next node freezes: no N2 and no A16 (TSMC next generation of manufacturing), no Rubin Ultra (Nvidia's next generation of chips). AI scaling plateaus for years.

TSMC

TRIGGERS

- Geopolitical event given geographic position
- Earthquake or power crisis
- CoWoS yield collapse
- Cyberattack

CONSEQUENCE

This is the most acute scenario. >90% of advanced chip production sits in Taiwan. Within weeks AI accelerator shipments collapse; auto, mobile and defence follow within a quarter. No fallback at scale to back up the demand.

NVIDIA

TRIGGERS

- CUDA misstep
- Architectural misstep
- Custom chip insurgency
- AI-capex slowdown

CONSEQUENCE

Substitutes exist (AMD, Google TPU, AWS Trainium) but all run on the same upstream pipeline. Most plausible outcome is a gradual market share erosion to 65-70%, not collapse. Software moat (CUDA, cuDNN, TensorRT) is the strongest defence.

A TSMC disruption is far more dangerous than ASML or NVIDIA.

Source: Company data, Morgans estimates

Bottlenecks worth watching

CoWoS packaging

Even at 140k wafers/month, demand still outruns supply. TSMC is offloading sub-steps to ASE and Amkor. Watch for yield surprises.

HBM memory

All three suppliers (SK hynix, Samsung, Micron) are sold out through 2026. The 16-layer HBM4 transition for Nvidia’s Rubin carries significant manufacturing challenges; any delay holds up the whole process.

High-NA EUV

TSMC deferred High-NA adoption to ~2029 on cost. Intel and Samsung are first movers. A surprise leader could shift competitive balance late-decade.

Power and water

Fabrication centres and AI data centres are consuming utility-scale power. Grid constraints in Taiwan, Arizona and Ireland are becoming a hard ceiling.

Geopolitical export controls

Each round of US-China chip controls reshapes the customer base. ASML cannot ship High-NA to China; TSMC's mainland revenue is shrinking.

Custom silicon ascendancy

Hyperscalers spending more on in-house chips (Google TPU, AWS Trainium, Microsoft Maia). Doesn't displace NVIDIA but caps its share.

Source: Company data, Morgans estimates

The key takeaways

What this means for the AI economy

1

The three are inseparable

Any view on AI exposure that names NVIDIA without also accounting for TSMC and ASML is incomplete.

2

The constraint is physical, not demand

AI revenue is gated by EUV machines, CoWoS lines and HBM stacks, not by customer.

3

Concentration risk is structural

One country (Taiwan), one company (ASML), one software stack (CUDA). There is no diversified path to current AI capability.

4

ASML is the most overlooked of the three

Lower public profile than NVIDIA, but a stricter monopoly.

5

The downside is not symmetric

A TSMC disruption is far more dangerous to global AI than a NVIDIA stumble.

The bottom line

NVIDIA designs the brain, TSMC builds it, ASML provides the tools without which it could not exist.

Each of these three companies is, today, irreplaceable.

The AI economy rests on a triangle with no diagonals, every side is load-bearing.

Source: Company data, Morgans estimates

Disclaimer

This presentation is provided for general information purposes only and is not intended as an offer to enter into any transaction. The information contained within this document is not necessarily complete and its accuracy can not be guaranteed. We have prepared this presentation without consideration of the investment objectives, financial situation or particular needs of any individual investor.

Before a client makes an investment decision, a client should, with or without Morgans' assistance, consider whether any advice contained in the presentation is appropriate in light of their particular investment needs, objectives and financial circumstances. It is unreasonable to rely on any recommendation without first having spoken to your adviser for a personal recommendation.

The information contained in this presentation has been taken from sources believed to be reliable. Morgans Financial Limited does not represent that the information is accurate or complete and it should not be relied on as such. Any opinions expressed reflect Morgans' judgment at this date and are subject to change. Morgans and/or its affiliated companies may make markets in the securities discussed. Further Morgans and/or its affiliated companies and/or their employees from time to time may hold shares, options, rights and/or warrants on any issue included in this presentation and may, as principal or agent, sell such securities.

The Directors of Morgans advise that they and persons associated with them may have an interest in the above securities and that they may earn brokerage, commissions, fees and other benefits and advantages, whether pecuniary or not and whether direct or indirect, in connection with the making of a recommendation or a dealing by a client in these securities, and which may reasonably be expected to be capable of having an influence in the making of any recommendation, and that some or all of our representatives may be remunerated wholly or partly by way of commission.

The presentation is proprietary to Morgans Financial Limited and may not be disclosed to any third party or used for any other purpose without the prior written consent of Morgans.

Morgans Financial Limited ABN 49 010 669 726 AFSL 235410
A Participant of ASX Group

Principal Office: Level 29, Riverside Centre, 123 Eagle Street, Brisbane QLD 4000

This document has been prepared by Morgans Financial Limited in accordance with its Australian Financial Services License AFSL no. 235410. The views expressed herein are solely the views of Morgans Financial Limited.

Stockbroking | Wealth Management | Corporate Advice