RESEARCH & INSIGHTS

THE HIDDEN LAYER SERIES | ARTICLE 9

By Joe Batista | November 2026 | 11-minute read

To subscribe to the series, complete the form

One Island

Advanced-chip production is spreading beyond Taiwan. The harder question is how much interruption those investments let a buyer withstand

The Hidden Layer series maps where value is created and captured across the AI stack, one chokepoint at a time. Articles 1 and 2 drew the map and showed how bottleneck power migrates through it. Articles 3 to 7 followed constraints on building AI: packaging capacity, the energized megawatt, the rules that ration it, the cost of money and code already written. Article 8 moved to the constraints on putting AI to work: who can prove an agent was allowed to act, and who will stand behind it. Article 9 returns to the physical floor of the stack: the factories where the most advanced chips are made, and where those factories are.

CORE INSIGHTS

1. **Diversification is real.** NVIDIA says its Blackwell chips reached volume production at TSMC's Arizona fab in October 2025, and TSMC plans $265 billion of US investment. These are changes, not just announcements.
2. **TSMC's manufacturing frontier is still concentrated.** TSMC's disclosed volume-production sites for N2, its newest process, are in Taiwan. Its second Arizona fab is scheduled to begin 3-nanometer volume production in the third quarter of 2027.
3. **The exposure is broader than any one node.** Current AI chips such as Microsoft's Maia 200 are made on TSMC's 3-nanometer process, and advanced packaging is a further dependency.
4. **Three prices, one not visible.** The cost of manufacturing elsewhere and the price of securing supply are both visible. We identified no comparable, publicly disclosed price for withstanding a prolonged interruption.
5. **A shock need not move the cluster.** After Thailand's 2011 floods, Western Digital restarted production where it stood, and its latest annual report still lists manufacturing there. Buyers should not assume a disruption will redistribute production.

The One-Island Paradox

The AI economy is being built across continents. TSMC's manufacturing frontier remains on one island.

The diversification is real. In October 2025, NVIDIA said its Blackwell chips had reached volume production at TSMC's fab in Arizona. TSMC says its planned Arizona investment now totals $265 billion, and the US government has awarded TSMC Arizona up to $6.6 billion in direct funding.

But geographic diversification and operational replacement are different achievements. On October 15, 2026, TSMC reported third-quarter revenue of $46.72 billion, up 41.2% from a year earlier, with N2, its newest process, at 8% of wafer revenue. Chairman and Chief Executive C.C. Wei said “N2 is ramping successfully in multiple phases at both Hsinchu and Kaohsiung sites,” both in Taiwan. TSMC’s second Arizona fab, for 3 nanometers, is scheduled to begin volume production in the third quarter of 2027. Demand, Wei said, “actually is still far above our capability today to supply.”

For an AI buyer, neither a fab count nor an investment total answers the continuity question. The relevant test is how long the business could meet its commitments if new accelerator deliveries stopped, and whether a tested alternative could take over in time.

This article separates three things often grouped under resilience: the cost of manufacturing elsewhere, the price of securing supply, and the ability to withstand a prolonged interruption. The evidence shows visible costs for the first two. It does not establish the third.

When we planned this series, the hidden-value story for this layer read: *nearly all leading-edge AI chips come from one company in one place, and the risk is priced as if it can’t happen.* We tested both halves. The first holds only in a narrower form. The second our evidence can neither confirm nor refute, and we explain why.

> *“The newest node still starts on one island. The money to change that is real, and so far it is buying the nodes behind it.”*

Where the Leading Edge Sits

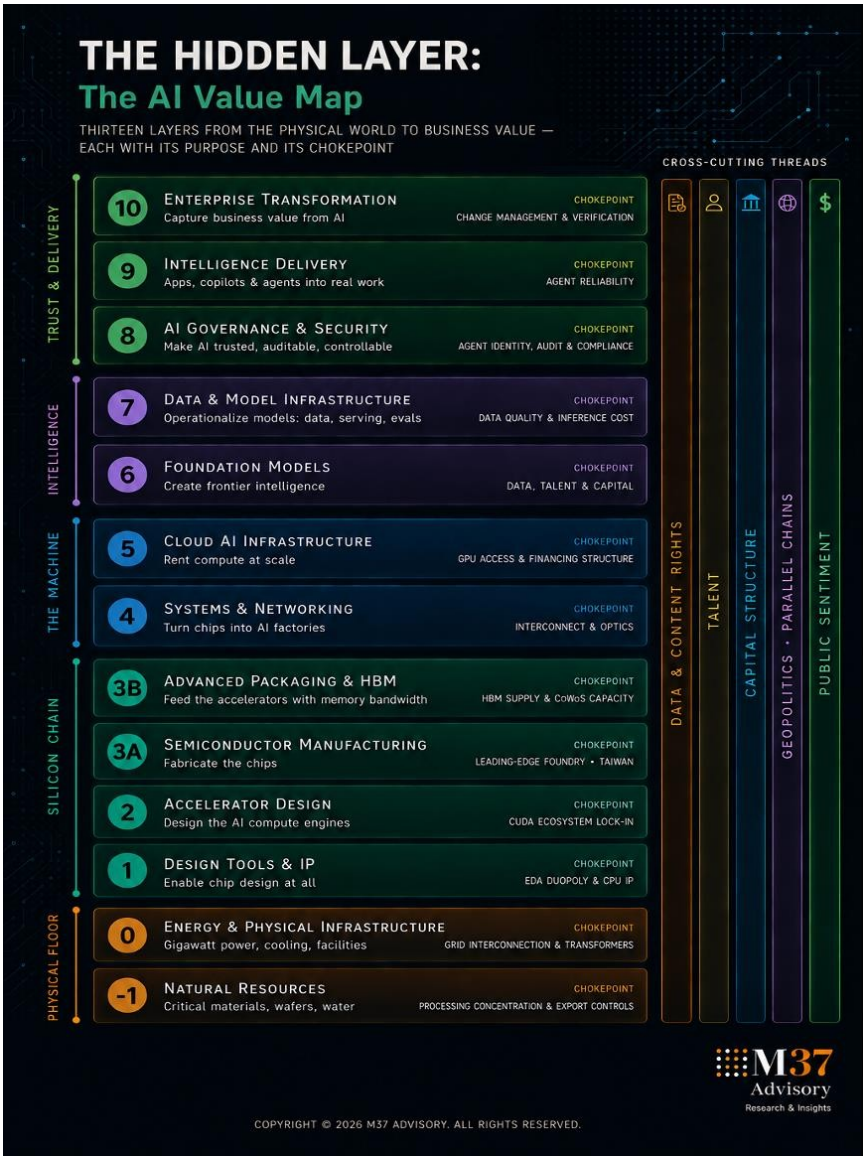

The Hidden Layer: The AI Value Map. 13 layers, each with its purpose and its chokepoint.

ON THE MAP: Layer 3A — Semiconductor Manufacturing | Chokepoint: Leading-Edge Foundry · Taiwan | This article: where TSMC's newest process is made, and how much continuity the money spent elsewhere buys

> *This article does not argue that all advanced chipmaking is in one place, or that diversification fails. It argues something narrower: TSMC's newest process starts in one place, and the money being spent to diversify buys capacity on the processes behind it. How much interruption a buyer could survive is a separate question, and the spending does not answer it.*
>
> — The thesis of this article

Two boundaries keep the argument honest. **First, "leading edge" here means TSMC's newest process in volume production,** not an industry ranking; Intel's 18A, a 2-nanometer-class process, is made in Arizona. **Second, the node is not the whole exposure.** N2 was 8% of TSMC's wafer revenue last quarter; the 3- and 5-nanometer processes were 33% and 28%. Current accelerators such as Microsoft's Maia 200 are built on 3 nanometers, and advanced packaging adds another dependency.

Two of the map's cross-cutting threads run through this layer: **Geopolitics · Parallel Chains** and **Capital Flows**. Article 2's model says bottleneck power grows with Constraint × Market Structure × Customer Urgency, enabled by Capital Capacity. We apply it below, with the evidence graded rather than scored.

Three Examples

① What Has Moved, and What Hasn't

By revenue, TrendForce estimates TSMC's share of the foundry market at 72.5% in the second quarter of 2026. That revenue-based measure spans process generations and customers; it does not measure TSMC's share of AI-accelerator production. The narrower argument rests on product disclosures and production locations.

By product. AMD relies on TSMC "for the production of all wafers for microprocessor and GPU products at 7 nanometer (nm) or smaller nodes." NVIDIA uses "foundries, such as" TSMC and Samsung "to produce our semiconductor wafers." Microsoft's Maia 200 is built on TSMC's 3-nanometer process.

By location. NVIDIA says Blackwell reached volume production at TSMC's Arizona fab in October 2025. TSMC's disclosed N2 volume-production sites are in Taiwan, and its second Arizona fab, for 3 nanometers, is scheduled for the third quarter of 2027.

Alternatives. Intel says its first 18A products reached high volume at the start of 2026, and that 18A-P, designed for future Intel products and external foundry customers, entered risk production in June 2026. NVIDIA says its Groq 3 LPX, a data-center AI inference accelerator that extends its Vera Rubin platform, entered full production in August 2026; in March, Chief Executive Jensen Huang said Samsung manufactures the chip.

Three different things can be called a "second source," and they should not be confused. An **alternative supplier** is another foundry producing a qualifying accelerator. A **second manufacturing source** is another qualified production route for a specific product. A **workload substitute** is another accelerator that can run the buyer's work acceptably. NVIDIA's Groq 3 LPX shows an alternative production route for one kind of accelerator at a non-TSMC foundry. It does not establish a second source for NVIDIA's TSMC-made GPUs, and NVIDIA offers it to work alongside Vera Rubin systems, not in place of them.

IN BRIEF: WHAT TO WATCH

- **A non-Taiwan date for TSMC's newest process.** When TSMC discloses a volume-production site outside Taiwan for its newest family, the gap becomes measurable.
- **A complete route.** The first named data-center accelerator made and packaged outside Taiwan at volume.
- **Outside customers in volume on Intel's 18A-P.**

> *"Two factories are not necessarily two sources. A second source is a chip you can actually make, and package, somewhere else."*

② What the Spending Buys

TSMC's first Arizona fab entered high-volume production at the end of 2024; its second, for 3 nanometers, is scheduled for the third quarter of 2027; construction of the third began in 2025. In July 2026 Wei announced

an additional $100 billion, saying “there’s probably additional four or more fabs will be built”; TSMC’s investor relations director put the total at $265 billion. No timeline was given. A commitment is not money spent, and money spent is not capacity installed.

Three parties are paying. **TSMC’s margin:** Chief Financial Officer Wendell Huang expects overseas fabs to reduce gross margin by about 2 percentage points in 2026, widening to 3 to 4 points in the next few years; that is a company-wide effect, not a wafer premium. **Customers:** asked about pricing, Wei said it “depends on the market segment, depends on the location, and depends on the technology,” and NVIDIA says that “to secure future supply and capacity, we have paid premiums, provided deposits.” **Governments:** TSMC Arizona was awarded up to $6.6 billion in CHIPS Act direct funding. Subsidies are industrial policy, not compensation for Taiwan-specific risk.

Mechanism	What it buys	Evidence of a price
Cost of geographic diversification	Capacity in a second location	Visible: margin dilution, location-dependent pricing, planned investment, subsidies
Price of supply assurance	A place in the allocation queue	Visible: premiums, deposits, prepaid capacity (NVIDIA)
Price of catastrophic exposure	Ability to withstand a prolonged loss of production	No comparable, publicly disclosed price identified in the materials reviewed

IN BRIEF: WHAT TO ASK YOUR SUPPLIERS

- **Where is my chip made, and where is it packaged?**
- **Which mechanism am I paying for?** A location premium buys capacity elsewhere; a deposit buys a place in the queue. Neither, by itself, buys continuity.

“The price of distance is visible. The price of priority is visible. The price of surviving a prolonged interruption is not.”

③ Who Carries the Tail?

Insurance. TSMC says its fabs and equipment are insured, “subject to some limitations,” against “earthquake, fire, typhoon, and other risks.” That cover has paid out: TSMC recognized about NT$3 billion of earthquake losses in the second quarter of 2024 and NT$5.3 billion in the first quarter of 2025, both net of insurance. The filing does not say whether war or a blockade is covered. A peril not listed in a public filing is not necessarily excluded, and a supplier’s insurance says nothing about what its customers can recover.

Disclosure. TSMC lists “the risk of outbreak of war or hostilities.” AMD says a TSMC shortfall “could have a material adverse effect.” NVIDIA refers to geopolitical tensions in regions where it relies on suppliers. We found no cited basis for a probability of conflict or blockade, and we offer none.

The filings acknowledge the risk but do not establish how a prolonged interruption’s losses would be allocated across the supply chain: among chip designers, cloud providers, their customers, insurers or governments. For a downstream buyer, that uncertainty is itself an exposure worth quantifying.

IN BRIEF: WHAT TO CHECK

- **Your insurance:** does business-interruption cover extend to a supplier's country, and what do its exclusions say?
- **Your contracts:** which force-majeure clauses would release your hardware or cloud supplier from delivering?

What This Changes for the Business

The planning question shifts from "Can we get enough AI capacity?" to "How much resilience have we actually bought?" Three things are often treated as one:

- **Capacity redundancy:** additional output that can serve demand.
- **Geographic redundancy:** production spread across locations.
- **Operational substitutability:** the demonstrated ability to move a chip design or a workload to an alternative without unacceptable time, cost or performance penalties.

Operational substitutability is the decisive test. Alternatives exist, but their ability to replace a disrupted leading-edge supply chain at the required scale, cost and speed remains unproven. A cloud provider can open a second region without a second semiconductor supply chain; a chip designer can secure capacity without a qualified alternative route.

The following business implications are M37 inferences, not reported findings.

- **For a technology CEO:** which customer commitments would you break if accelerators arrived two quarters late, and does your board know which ones?
- **For a chief revenue officer:** how much of next year's pipeline depends on AI capacity going live on a hardware delivery date?
- **For a channel founder:** the opening is in substitutability: mapping where customers' chips are made and packaged, and designing fallbacks that have been tested.
- **For a board member:** when management reports that the company has diversified, ask whether it has bought redundancy or substitutability, and how many weeks of interruption the business could absorb.

"Spending on resilience is not the same as having it."

The Pattern: The Flood That Didn't Move the Factories

In the quarter ended July 1, 2011, Western Digital shipped about 54 million hard drives, roughly 60% from its sites in Thailand. In October, flood water breached the defenses of the Bang Pa-in industrial park and submerged its main manufacturing site. WD estimated that December-quarter industry shipments would be limited to about 120 million drives, against demand of 170 million to 180 million. It shipped 28.5 million drives that quarter, against 52.2 million a year earlier.

In its fiscal 2012 annual report, WD recorded $214 million of net flood-related charges, after $21 million of insurance recoveries and reimbursements received by then, and said final losses could exceed its insurance limits. It restarted drive production in Bang Pa-in on November 30, 2011, and announced a new slider fab in Penang, Malaysia, a second source for one component. Its fiscal 2025 annual report still lists manufacturing at Bang Pa-In and Prachinburi, Thailand.

Our interpretation: the lesson is not that diversification fails. It is that rebuilding the incumbent cluster may remain attractive after a shock, while targeted changes, such as a second site for one component, happen around it. Buyers need continuity plans that do not assume a disruption will permanently redistribute production. The precedent has limits: a flood recedes, a geopolitical disruption could last longer, and hard drives in 2011 were a mature product.

> *"Concentration can survive a catastrophe because the economics that created it survive too. The question is whether the buyer can survive the next interruption."*

The M37 Lens: Article 2's Model, Applied to One Island

Article 2 asks four questions of any bottleneck: Is there a shortage? Can suppliers exploit it? Can buyers wait? Can buyers pay? Here we grade the evidence for each dimension rather than present numerical scores; our band judgments, which are M37 judgments, are in Appendix B.

Dimension	Established by available evidence	Indicated, but incompletely measured	Not established from public sources
Supply headroom	TSMC: demand "far above" supply (Oct 2026)	—	Duration of tightness; spot versus contract pricing
Supplier concentration	TSMC's newest process made at disclosed Taiwan sites; AMD's ≤7nm wafers all at TSMC	High concentration for AI accelerators (foundry revenue share is a proxy)	Share of AI-accelerator output by foundry
Qualification time	—	—	Time to move a leading-edge accelerator between foundries
Replication cost	Multibillion-dollar scale of geographic replication	—	Cost of reproducing a particular process or facility
Switching cost	—	—	Cost of porting a leading-edge design
ROI urgency	Large 2026 capital commitments (Amazon ~$220B cash capex; Meta $135–145B incl. finance leases)	Urgency to deploy	Whether buyers can defer purchases
Capital availability	TSMC 2026 capital budget $60–64B	—	Buyers' financing capacity

M37 VIEW: Layer 3A holds two bottlenecks. A **demonstrated capacity bottleneck**: demand exceeds supply, and supply is allocated, as Article 3 found one layer up. And a **latent geographic bottleneck**: TSMC's newest process sits in one place, and the event that would make that binding has not occurred. TSMC's position is strong; demonstrated pricing power is not shown, because its 66.4% gross margin is company-wide, not a wafer price. That placement is our judgment.

WHAT WE MEAN BY "ONE ISLAND"

Observed: TSMC's disclosed volume-production sites for its newest process are in Taiwan; diversification on older processes is under way; the costs of manufacturing elsewhere and of securing supply are visible; the filings do not show how a prolonged interruption's losses would be allocated.

Inferred, the concentration thesis: our base case is that TSMC's newest process family will remain Taiwan-only through June 2028. This is an inference from disclosed production plans, not company guidance.

Hypothesized, the "margin, not tail" hypothesis (unproven): buyers and suppliers are paying for capacity elsewhere and for priority, without yet buying a demonstrated substitute route for a prolonged loss of Taiwan output.

The One-Island Test: Four Questions You Should Be Asking

1. **Dependence:** what share of your AI compute spend, owned or rented, runs on chips made or packaged in one place?
2. **Runway:** how many weeks could you meet your commitments if new accelerator deliveries stopped?
3. **Allocation of loss:** for each hardware contract, cloud agreement and insurance policy, does a loss of supply fall on you, your supplier or an insurer?
4. **Substitutability:** for your three most important AI workloads, how many months and dollars would it take to run them on a different supply route, and has anyone tested it?

The Predictions

Same rules as Articles 2 to 8: checkable by anyone, including where we were wrong. Counting windows run from publication to June 30, 2028; we report in October 2028. These tests measure observable signals; none can measure how catastrophic risk is priced or transferred. Definitions, baselines and counting rules are in Appendix A.

Prediction	Supports	Weakens	Tests
The newest node starts at home	On June 30, 2028, every volume-production site TSMC discloses for its newest process family is in Taiwan	TSMC discloses volume production of that family at a site outside Taiwan	Thesis
No new alternative production route at scale	During the window, none of the seven named designers announces a data-center AI accelerator in volume production at a foundry other than TSMC (NVIDIA's Groq 3 LPX, announced before the window, is the baseline and does not count; a later product does). Exactly one: reported as early diversification that partly weakens the thesis	Two or more do	Thesis

Prediction	Supports	Weakens	Tests
Geographic premiums without demonstrated substitutability	Both: (a) TSMC or a named designer says US-made wafers or chips cost more than Taiwan-made ones; (b) during the window, no company discloses a named data-center accelerator made and packaged outside Taiwan at volume	Such an end-to-end route is disclosed. If (b) holds without (a): untestable	Hypothesis (indicators only)
Arizona's second fab	Volume production by Sept 30, 2027; Oct 1, 2027 to Jun 30, 2028 reported as late but completed	None by June 30, 2028	Supporting
Diversification gets costlier	TSMC reports overseas dilution of at least 3 percentage points in any full year through 2027	Below 3 points every year; reported either way	Supporting

The Bottom Line

The concentration is measurable, but not with a single statistic. Foundry revenue share, process location, product manufacturing and packaging dependencies describe different parts of the exposure. Diversification is also real. What investment announcements do not establish is how much disrupted output could be replaced, how quickly, and for which customers.

On the series' original claim, our review does not establish whether markets adequately price a prolonged Taiwan disruption. It does establish that geographic diversification and supply assurance carry visible costs. Those costs do not tell a buyer how much interruption it could withstand. Until a second place can replace the first, diversification reduces exposure; it is not yet an escape route.

The practical test is therefore not how much the industry has committed to spend. It is how many weeks your business could continue operating, which obligations would fail first, and whether an alternative supply route or workload platform has been tested.

> *"Money committed is evidence of a plan. Production delivered is evidence of capacity. Continuity demonstrated is evidence of resilience."*

How This Article Was Tested

Sources. Facts come from primary documents wherever they exist: TSMC's release, call transcripts and 2025 annual report; annual and quarterly reports of AMD, NVIDIA, Intel, Amazon, Meta and Western Digital; company releases and blogs; and US Commerce Department documents. Where a fact rests on reporting, the outlet is named. Investment and capacity plans are reported as plans, not capacity. Every statistic and quotation is listed, with its original source, in a separate verification record.

Three levels of claim. Observations, inferences and the hypothesis are labeled separately. Both halves of the original hidden-value story were tested as claims; the first was narrowed, and the second could not be settled with the evidence available.

What we left out. Our rule: if a claim could not be verified at its original source, it was removed. That excluded figures that Taiwan makes "over 90%" of advanced chips; TSMC's reported target for the Arizona

share of its 2nm-and-beyond capacity, which we could not check against the official transcript; reports of where Arizona-made Blackwell wafers are packaged; the terms of US Section 232 chip duties; foundry attributions for chip designers that do not name their foundry; reported foundry contracts not confirmed by both parties; a count of TSMC fabs in Taiwan; reported wafer price increases; a 2011 per-drive price not in WD's filed release; and any probability of conflict.

Disclosure. None of the companies named in this article is a current or former M37 Advisory client.

Predictions. Thresholds, the named list, windows, baselines and counting rules were fixed before publication and will not be changed.

About the Author

Joe Batista advises boards, C-suite executives, and go-to-market teams on technology value creation, growth strategies, and the opportunities hidden beneath major technology shifts. He has spent more than 40 years in the technology industry, guiding companies on growth and technology decisions. His career spans the minicomputer era at Digital Equipment Corporation (DEC) through client-server, cloud, virtualization and now AI. He has held executive roles at Dell Technologies and Hewlett Packard Enterprise.

Through M37 Advisory, Joe helps leadership teams recognize recurring patterns, uncover second-order effects, find the opportunities others overlook, and identify the strategy worth betting on. His work is grounded in a simple belief: What if the biggest obstacle to your growth isn't the competition, but the way you think?

> *"Most executives have been bombarded with 'theories' about what might work. Based on our research, Joe's extensive playbook is rooted in creating alignment of interest and business outcomes that leaders should take notice of."*
>
> — Jim Dickie, Research Fellow & Sales Mastery Founder, CSO Insights

Email: Chief.Creatologist@M37-Advisory.com | Website: www.m37-advisory.com | Idea Vault: www.m37-advisory.com/idea-vault

NEXT IN THE SERIES

Article 10: "The Wire Between the Chips." We move up the map to Layer 4, Systems & Networking, where the chokepoint is interconnect and optics. Every chip in this article is built to work alongside thousands of others, and as clusters grow, moving data between chips becomes the limit. Performance is set by the slowest link.

To subscribe to the series, complete the form.

Appendix A: Prediction Methodology

- **Sources:** official company or government documents; where none exist, two independent major news reports.
- **Volume production:** the company says “volume production,” “high-volume manufacturing” or “mass production.” Risk production, pilot lines and sampling do not count.
- **Newest process family:** the most advanced process TSMC reports in volume production on June 30, 2028. N2 and its variants are one family; A16 and A14 are each their own. If a newer family becomes the newest before that date, the test applies to it. Disclosed sites are those TSMC names; we do not treat an unnamed site as evidence either way.
- **Named designers:** NVIDIA, AMD, Broadcom, Alphabet, Amazon, Microsoft and Meta. Subsidiaries count as their parent; a chip designed with a partner counts once, for the company that deploys it.
- **Data-center AI accelerator:** a GPU or custom accelerator deployed for AI training or inference in data centers. The compute die counts; memory, base dies and packaging do not, except in the end-to-end test.
- **Alternative production route (test 2):** another foundry produces a qualifying accelerator in volume, as announced during the window. Announcements before publication, such as NVIDIA’s Groq 3 LPX in August 2026, are the baseline and do not count. This is not the same as a second source for an existing TSMC-made product.
- **End-to-end route (test 3):** a specific named product (not an architecture or product family) whose wafer fabrication and advanced packaging both take place outside Taiwan in volume, as disclosed by the designer, TSMC or the packaging supplier. Only a disclosure first made during the window counts; wafer fabrication alone, including NVIDIA’s October 2025 Blackwell announcement, does not.
- **Location premium:** an explicit statement that US-made wafers or chips cost more than equivalent Taiwan-made ones. A location premium is a price for manufacturing in a second place, not by itself a premium for geopolitical risk.
- **First exception:** if exactly one named designer qualifies under test 2, the test reads as supported under the fixed threshold, and we report the exception as early diversification that partly weakens the thesis.
- **Expansions:** announced fabs, pledges and capacity plans are not production.
- **Insufficient evidence:** if no official text or two independent reports settle an item, it does not count, and we say so. A test with too little qualifying evidence is reported as untestable.
- **Judgments:** thresholds, the named list and the definitions are M37 judgments, fixed on publication.

Appendix B: M37 Band Judgments (Article 2 Rubric)

These are M37 judgments, not measurements. They place each dimension in Article 2’s low, medium or high band only where evidence allows; they are not combined into a single score.

Dimension	M37 band judgment	Basis and limit
Supply headroom	Medium (6)	Demand above supply stated (Oct 2026); duration and pricing not measured
Supplier concentration	High (9), on proxy evidence	Foundry revenue share and disclosed sites; not an AI-output share
Qualification time	Not scored	No primary source
Replication cost	High (10)	Program scale far above the $2B band; no per-facility cost isolated
Switching cost	Not scored	No primary source

Dimension	M37 band judgment	Basis and limit
ROI urgency	High, qualified (7)	Commitment shown; ability to defer not tested
Capital availability	Abundant (8)	Supplier capital budget; buyers' financing not assessed

Sources

PRIMARY DOCUMENTS

[1] TSMC. "TSMC Reports Third Quarter EPS of NT$29.43." October 15, 2026. Link

[2] TSMC. Third Quarter 2026 Earnings Conference Call, edited transcript (LSEG StreetEvents). October 15, 2026. Link

[3] TSMC. Second Quarter 2026 Earnings Conference Call, edited transcript. July 16, 2026. Link

[4] TSMC. Annual Report on Form 20-F for 2025. Filed April 16, 2026. Link

[5] TSMC. "TSMC January 2025 Revenue Report and the Statement on the Impact of Earthquake." February 10, 2025. Link

[6] Advanced Micro Devices. Form 10-K for fiscal 2025. Link

[7] NVIDIA. Form 10-K for fiscal 2026 (year ended January 25, 2026). Link

[8] NVIDIA. "The Engines of American-Made Intelligence: NVIDIA and TSMC Celebrate First NVIDIA Blackwell Wafer Produced in the US." NVIDIA Blog, October 17, 2025. Link

[9] Intel. Form 10-Q for the quarter ended September 26, 2026. Link

[10] Intel. "Intel Unveils Panther Lake Architecture: First AI PC Platform Built on 18A." October 9, 2025. Link

[11] Microsoft (Scott Guthrie). "Maia 200: The AI accelerator built for inference." Official Microsoft Blog, January 26, 2026. Link

[12] US Department of Commerce. "Biden-Harris Administration Announces CHIPS Incentives Award with TSMC Arizona." November 15, 2024. Link

[13] Amazon.com. "Amazon.com Announces Third Quarter Results." Form 8-K Exhibit 99.1, October 29, 2026. Link

[14] Meta Platforms. "Meta Reports Third Quarter 2026 Results." Form 8-K Exhibit 99.1, October 21, 2026. Link

[15] Western Digital. "WD Reports Impact on HDD Production Due to Severe Flooding in Thailand." Form 8-K, October 12, 2011. Link

[16] Western Digital. "WD Provides Update on Effect of Thailand Floods." Form 8-K Exhibit 99.1, October 17, 2011. Link

[17] Western Digital. "WD Updates December Quarter Outlook, Reports on Recovery Effort in Thailand." December 1, 2011. Link

[18] Western Digital. "WD Announces Q2 Results and Operations Update." Form 8-K Exhibit 99.1, January 23, 2012. Link

[19] Western Digital. Form 10-K for fiscal 2012. Link

[20] Western Digital. Form 10-K for fiscal 2025, Item 2, Properties. Link

[21] NVIDIA. "NVIDIA Groq 3 LPX Now in Full Production With World-Class Speed for Agentic AI." August 24, 2026. Link

[22] NVIDIA. "Transcript: NVIDIA GTC 2026 Keynote with Jensen Huang," section 1:32:52. March 16, 2026. Link

[23] M37 Advisory. Articles 1, 2 and 3 of The Hidden Layer Series. 2026.

REPORTING AND COMMENTARY

[24] TrendForce. "Global Foundry Revenue Approaches US$53.49 Billion in 2Q26 as SMIC Narrows Market Share Gap with Samsung." September 9, 2026. Link

© 2026 M37 Advisory. All rights reserved.