LTM

WHITEPAPER

Two Clocks, One Balance Sheet

Why the ability to keep options open is quietly becoming the chip industry's real superpower



Table of Contents

Executive Summary 5

Introduction 5

The Problem Beneath the Boom 6

Implications 9

What This Means For Leaders 10

What's The Solution 11

The Bottom Line 12

References 13

Authors 14

Executive Summary

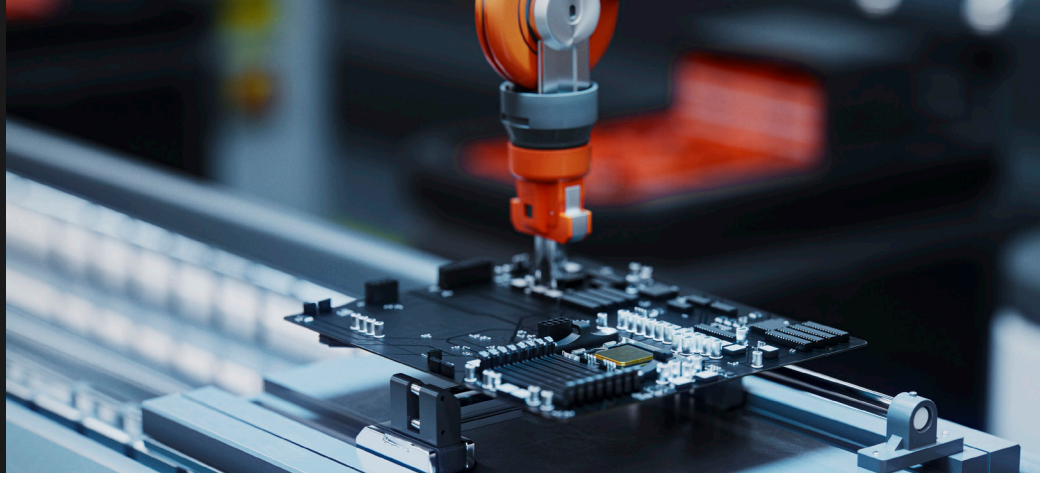

A global semiconductor manufacturer recently faced a problem that many industry leaders will recognize. Market volatility was putting pressure on margins, inventory, cash, and day-to-day operations. At the same time, the business could not afford to slow investment in technologies that would matter in the next cycle. Information existed across the organization, but demand, supply, capacity, operational, and financial signals sat in different places. Leaders could see each part, but understanding the knock-on effect of a change was harder. LTM brought these signals into a connected, trusted enterprise view. Teams could spot issues sooner, work from the same evidence, and make decisions with less disruption to the wider business.

The work reflected the tension at the heart of this paper. The quarters clock called for control, speed, and operational discipline. The years clock still required commitment to future technology and growth. The company did not solve the problem by choosing one horizon over the other. Better-connected data and operations gave leaders a practical way to manage both. They could respond to immediate signals, judge the broader consequences, and adjust course while keeping long-term priorities intact. In effect, LTM helped build an operating foundation that supported the next quarter without closing off options for the next decade.

The two clocks need not pull the business in opposite directions. A stronger operating foundation can protect current performance and keep long-term options open.



Figure 1. One operating foundation connects immediate performance with long-term capability.

Introduction

Talk to semiconductor executives in 2026, and the contraction is hard to miss. On the surface, the industry rarely looks stronger. AI silicon is reshaping growth, while hyperscalers and governments are committing unprecedented capital. Yet the decisions facing leadership teams have become more difficult, not less. Chips now sit at the center of AI infrastructure, national security, and digital sovereignty. As a result, an investment decision that once looked mainly commercial can now carry strategic and geopolitical consequences.

The Problem Beneath the Boom

Rarely stronger on paper, seldom harder to steer

The tailwinds are substantial. Gartner expects semiconductor revenue to exceed USD 1.3 trillion in 2026, a 64% increase and the fastest growth in two decades. AI silicon is expected to contribute nearly 30% of that total. NVIDIA, at the center of the accelerator market, ended fiscal 2026 with revenue of USD 215.9 billion, up 65% year on year, led overwhelmingly by its data-center business.

Hyperscalers are raising AI infrastructure budgets by more than 50%, and policymakers are reinforcing the momentum. Washington's CHIPS program has committed direct grants, including USD 6.6 billion to TSMC's Arizona operations and USD 7.86 billion to Intel. Europe's Chips Act has mobilized more than EUR 52 billion, while Japan has committed roughly USD 15 billion to Rapidus's drive toward 2-nanometer production.

Why, then, does an industry with such strong numbers feel so difficult to lead? The answer is not simply volatility. The role of the chip itself has changed. Semiconductors are no longer only components inside finished products; they are infrastructure for AI, security, and digital sovereignty. That change raises the stakes of every capital decision. Behind the headline growth, leadership teams are being pulled by two pressures that do not move in step.

The first pressure: Demand that can no longer be confidently predicted

For years, semiconductor companies could rely on demand forecasts with enough confidence to support a fab investment. That confidence has weakened. Demand can move quickly, the impact can be severe, and the trigger may come from customers, policy, supply constraints, or a technology shift. Four forces now deserve particular attention.

01

Leaner Inventory

Customers order closer to real-time need, reducing the buffer against volatility.

02

Shortage And Glut

Leading-edge capacity tightens while selected mature segments remain oversupplied.

03

Policy Volatility

Trade rules and export controls can reshape market access and investment assumptions.

04

Lopsided Growth

AI-linked segments accelerate while parts of the automotive, industrial, and consumer markets recover unevenly.

Figure 2. Four forces are shortening the useful life of a single-point forecast

- **Customers are carrying less inventory:** After the 2023-24 correction, many buyers moved closer to just-in-time purchasing. That leaves the value chain with less stock to absorb a sudden change in demand.
- **Shortage and oversupply now coexist:** TSMC's advanced nodes and Chip-on-Wafer-on-Subsrate (CoWoS) capacity are booked through at least 2027, while selected mature segments remain oversupplied. CEO C.C. Wei called AI demand "insane" and warned that supply could trail demand "for years."
- **Trade policy has become a moving variable:** In January 2026, the United States imposed a 25% Section 232 tariff on imported chips, while Section 301 duties on Chinese semiconductors rose to 50%. Rules affecting NVIDIA's H20 and H200 changed several times within 12 months. The expiry of Validated End User status for TSMC's, Samsung's, and SK hynix's China fabs added another annual decision point.
- **Growth is uneven across the portfolio:** Advanced compute, high-bandwidth memory, and sophisticated packaging are accelerating with AI infrastructure. Automotive, industrial, and parts of the consumer electronics markets are recovering more slowly. WSTS expects discrete semiconductors to contract in 2026, in part because of soft automotive demand.

Growth is real - but lopsided: AI-linked segments race ahead

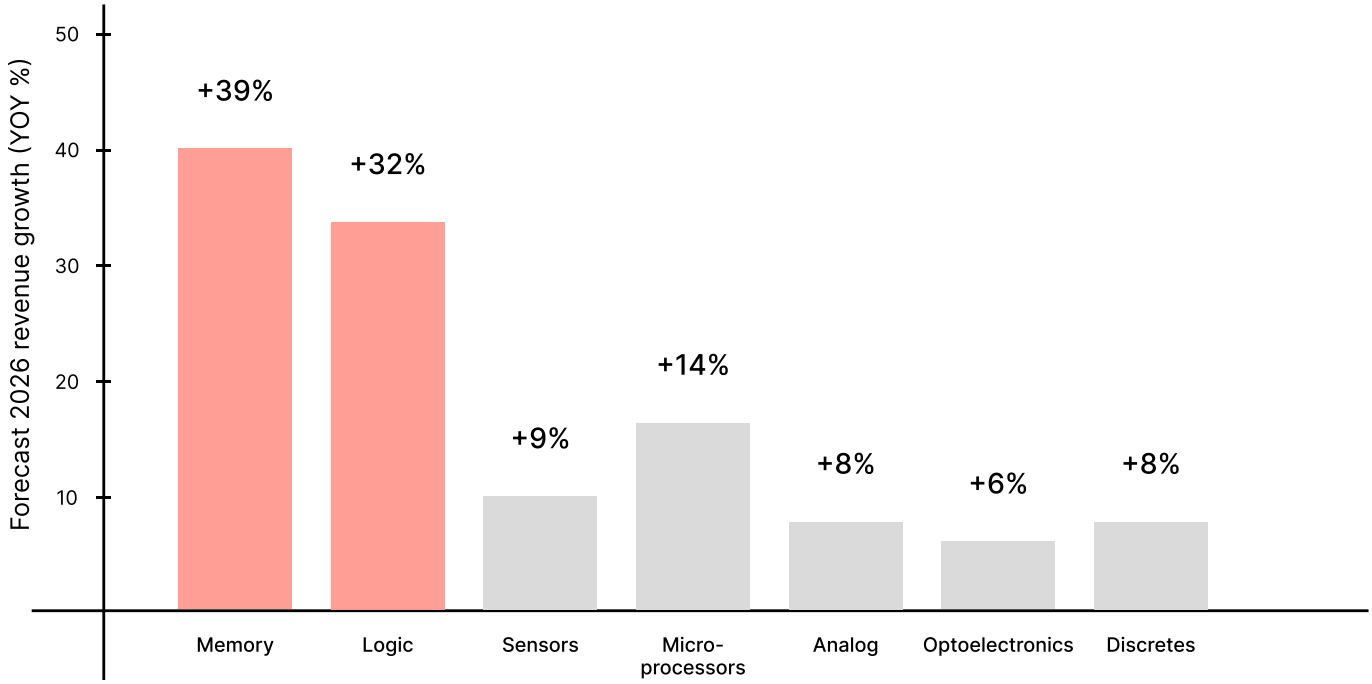

Source: WSTS Autumn 2025 forecast. Memory and logic — the AI backbone — lead; discrete and analog lag

Together, these shifts change the management question. The priority is no longer to produce a perfect forecast. It is to understand how quickly the business can respond when the forecast proves wrong.

The second pressure: The need to keep spending, even when instinct says hold back

This creates difficult capital choice. Volatile demand encourages caution, but the next technology cycle still requires investment. Leaders must protect margins during softer periods without underfunding the capabilities that will matter later. Both demands ultimately compete for the same balance sheet.

- The cost of remaining competitive continues to rise: TSMC has moved its 2026 capital spending toward USD 60 billion. Hyperscaler capital intensity has reached 45%–57% of revenue, a level traditionally associated with utilities.
- The investment agenda is also expanding: Silicon carbide, gallium nitride, electrification, and autonomy require funding before demand is fully visible. The combined market for SiC and GaN power devices is expected to exceed USD 14 billion within five years, with AI data centers adding a new source of demand.
- Customers want more computing performance and lower power consumption at the same time: Delivering both requires sustained engineering effort and capital. Neither improvement arrives quickly or cheaply.
- Mature platforms cannot simply be left behind: Companies must keep legacy nodes productive and profitable while funding leading-edge capacity. The balance shifts over time, but the work never fully stops.

History offers little comfort to companies that wait for certainty. Industry leaders often invest before demand becomes obvious. Holding back may support the current quarter, but it can weaken the company's position when the next cycle takes shape.

This is the practical meaning of the two clocks. The first is measured in quarters: earnings, inventory, and shareholder experience. The second is measured in years: roadmaps, fabs, and the slow work of technical progress. They seldom align, but major decisions must account for both.

> Markets can change within quarters. Building real capability takes years. Leadership now sits between those time frames.

Implications

When a single decision moves the whole board

The present cycle is unusually concentrated. Earlier upcycles were spread across PCs, phones, and consumer devices, with millions of buyers sharing the risk. This cycle depends heavily on a small group of hyperscalers. Their combined 2026 capital expenditure is expected to approach USD 725 billion, around 77% higher in one year, based on The Financial Times' tally from first-quarter earnings. A budget change by one cloud provider can therefore affect plans across the value chain.

The chain reaction is visible. Accelerator orders rise, foundries become constrained, memory producers reallocate high-bandwidth capacity, advanced packaging tightens and power utilities are pulled into the investment plan. A decision made by one company can alter assumptions for dozens of others. The sector now behaves less like a conventional cycle and more like a network. The same connection works in reverse: a packaging delay can hold back shipments, pause enterprise AI programs, and force software roadmaps to change.

A closer look reveals another complication. Many semiconductor companies are managing two different economies within same organization. AI-related businesses are expanding rapidly, while mature-node, automotive, industrial, and parts of consumer demand remain uneven. Treating the enterprise as though every business is moving in the same direction can hide the real capital and operating choices.

One Industry, Two Economies Under One Roof

Economy One

The Traditional Economy

Mature-node manufacturing, industrial automation, parts of consumer & auto — recovery remains uneven.

- Inventory normalisation continues
- Discrete chips set to decline in 2025 (WSTS)
- Automotive demand varies by geography
- Customers stay lean and selective

Economy Two

The AI Economy

Accelerators, HBM, optical interconnect, advanced packaging — growing at exceptional rates.

- AI silicon ≈ 30% of all chip revenue (2026)
- Nvidia FY26 revenue $215.9B, up 65%
- Hyperscaler AI capex up >50% YoY
- Advanced nodes & CoWoS sold out to 2027

The split leaders now manage: an AI economy running hot, a traditional economy still recovering

What This Means For Leaders

Not steering a ship. Managing a fleet.

Capital decisions now span businesses facing very different market conditions. Leaders must decide where to add capacity, which nodes to support, and when a fab commitment is justified. AMD CEO Lisa Su has described AI chip demand as “insatiable” and the market as being in the “early innings of a multiyear adoption cycle.” During the same period, SiC suppliers exposed to automotive demand dealt with price pressure and thin order books. These conditions can exist in the same industry, and sometimes inside the same company.

> **“Insatiable” demand, yet still the “early innings of a multiyear adoption cycle.”**
>
> Lisa Su, Chair and CEO, AMD

The boardroom question is no longer whether the market is growing. Leaders need to know which parts are growing, what is driving them, and how quickly the pattern could change. Operational efficiency still matters, but it cannot answer those questions on its own. Companies also need visibility beyond their internal operations, including suppliers, partners, and other dependencies, so that risk and business continuity decisions reflect the wider network.

> The next decade will not be won by budget size alone. It will favor leaders who can read the situation clearly and make right trade-offs.

What's the Solution

Stop forecasting harder. Start deciding better.

When uncertainty rises, companies often respond by adding data and building more sophisticated forecasts. Those tools help, but they do not remove the central problem.

AI workloads can shift faster than roadmaps, export rules may change during the year, and a customer can revise priorities long before a fab can be retooled. No single forecast remains reliable for the full life of a multibillion-dollar asset.

The practical task is to improve the way capital decisions are made, rather than trying to predict every turn.

Four changes can make the approach repeatable.

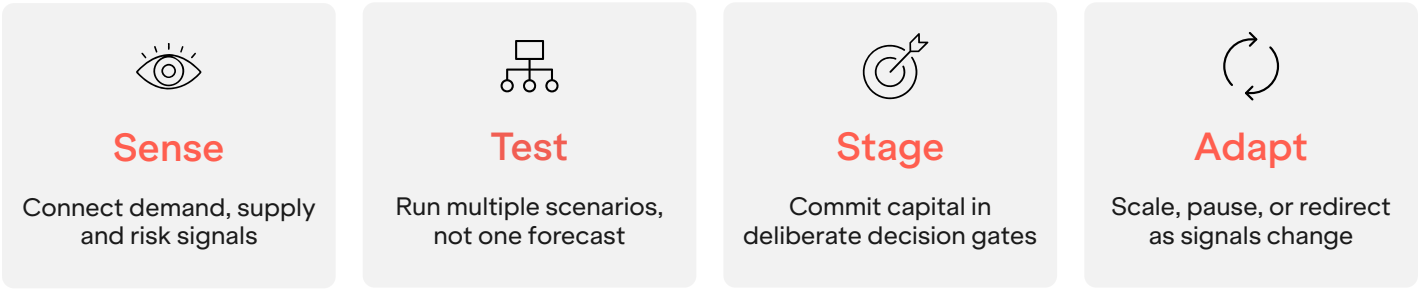

Figure 3. Optionality becomes repeatable when sensing, scenarios, capital, and execution work as one loop

- **Dynamic capital allocation:** Move beyond a fixed annual budget. Stage major commitments against explicit demand and risk signals, with agreed points to expand, pause, or redirect spending.
- **Scenario planning as standard:** Carry several demand cases in parallel such as AI-heavy, broad recovery, and prolonged automotive weakness, and stress-test every major investment against each, rather than wagering on a single projection.
- **AI-enabled sensing:** Use digital twins, machine learning yield tools, and control towers to connect demand, supply, and geopolitical indicators. The objective is not more dashboards; it is a shorter path from an important signal to a decision.
- **Optionality by design:** Build enough flexibility to serve more than one future without repeatedly redesigning the business. Geographic diversification, including TSMC's footprint across Taiwan, Arizona, and Japan and Intel's external foundry model, can create additional room to respond.

Technology can improve the quality and speed of these decisions, but leadership judgment remains essential. A model cannot fully price geopolitical risk, and a digital twin cannot decide whether another USD 20 billion of capacity is justified. Strong tools work best when paired with clear governance, operating experience, and a culture that revisits assumptions early.

> When assumptions change faster than the assets built around them, optionality becomes part of operating efficiency.

The Bottom Line

The semiconductor industry offers exceptional growth, but it also asks leaders to commit capital under persistent uncertainty. The companies that shape the next decade will not be distinguished only by larger budgets or more accurate forecasts. They will be able to act, review the evidence, and change direction before a weak bet becomes an expensive constraint. That is optionality in practical terms: commit where conviction is justified, preserve alternatives where uncertainty remains, and stop funding choices that no longer hold.

The work with the global semiconductor manufacturer illustrates how a connected operating foundation can support this discipline. By bringing critical signals into a more trusted enterprise view, LTM helped teams respond to volatility while keeping longer-term priorities in sight.

> In an industry governed by two clocks, advantage will come from making sound commitments while keeping enough room to adapt. Waiting for complete certainty is no longer a viable strategy.

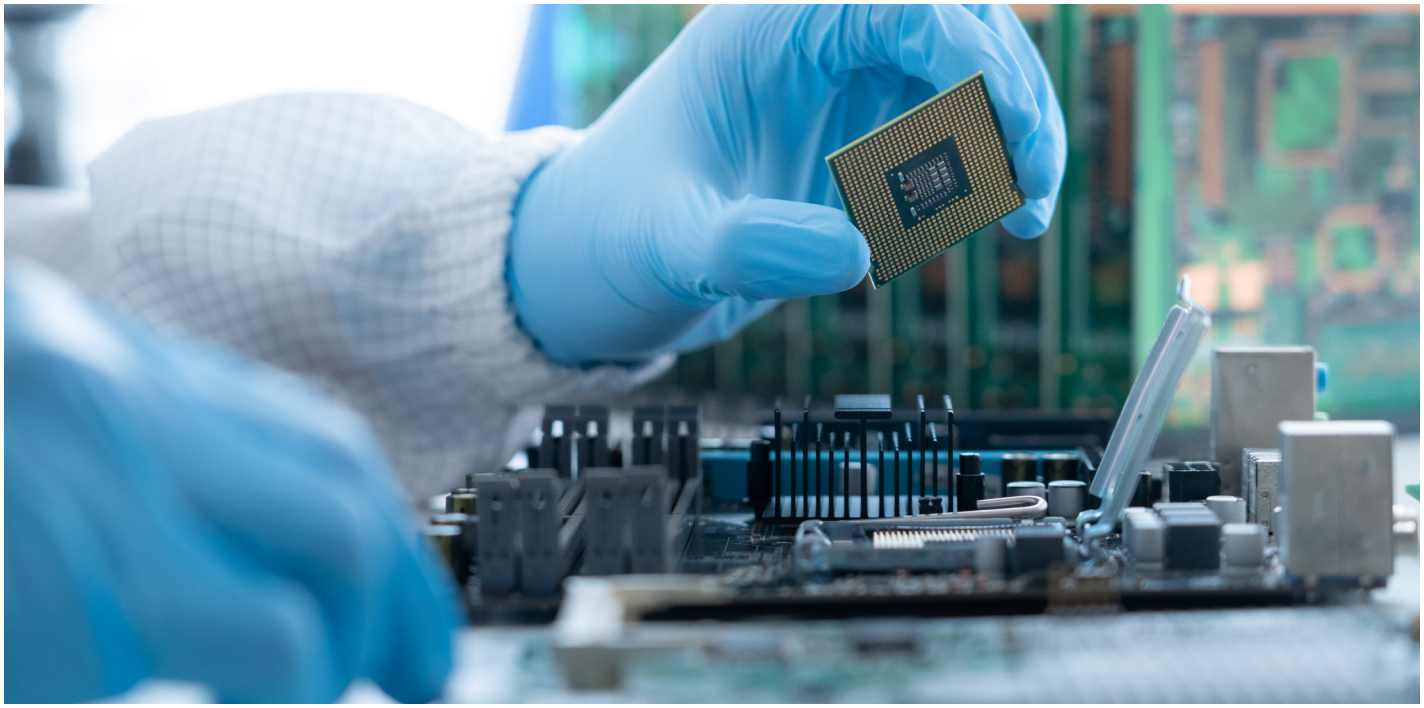

References

All figures and quotations above are drawn directly from the primary and industry sources listed below.

1. Gartner — Worldwide semiconductor revenue to exceed $1.3 trillion in 2026; AI ≈ 30% of revenue; hyperscaler AI spend up >50%: https://www.gartner.com/en/newsroom/press-releases/2026-04-08-gartner-forecasts-worldwide-semiconductor-revenue-to-exceed-us-dollars-one-point-3-trillion-in-2026

2. NVIDIA — Q4 & Fiscal 2026 results ($215.9B revenue, +65%; Huang quote on agentic AI inflection): https://investor.nvidia.com/news/press-release-details/2026/NVIDIA-Announces-Financial-Results-for-Fourth-Quarter-and-Fiscal-2026/

3. WSTS / ESIA — Autumn 2025 forecast: segment growth; discretes to decline on automotive weakness: https://www.eusemiconductors.eu/sites/default/files/ESIA_WSTS_PR_AuFc2025.pdf

4. Reuters via ECM — TSMC CEO C.C. Wei on "insane" AI demand, supply lagging "for years," pricing power (AGM, June 2026): https://ecmsource.com/tsmc-cc-wei-agm-ai-chip-pricing-power-june-2026/

5. Yahoo Finance — TSMC 2026 capex raised to $60–$64B; HPC 66% of revenue; Wei on "extremely robust" AI demand: https://finance.yahoo.com/technology/ai/articles/tsmc-revenue-surges-ai-chip-130445040.html

6. CNBC — AMD CEO Lisa Su on "insatiable" AI demand and multiyear adoption cycle: https://www.cnbc.com/2025/11/11/amd-lisa-su-growth-ai-analyst-day.html

7. CRN — AMD Q2 2026: Su on "early innings of a multiyear AI adoption cycle": https://www.crn.com/news/ai/2026/amd-q2-earnings-ceo-su-boosts-outlook-as-accelerator-cpu-sales-climb

8. CreditSights / ValueAdd — 2026 hyperscaler capex ~$725B, up ~77% YoY; capital intensity 45–57%: https://valueaddvc.com/blog/ai-hyperscaler-capex-compared-why-microsoft-google-meta-and-amazon-are-all-spending-at-once

9. AI Weekly - Amazon, Microsoft, Alphabet and Meta collectively planning around $725B of capital expenditure in 2026. That is up 77% from last year's record of roughly $410B: https://aiweekly.co/alerts/amazon-microsoft-alphabet-meta-plan-725b-ai-capex-in-2026

10. The White House — Proclamation: 25% Section 232 tariff on semiconductors (Jan 2026): https://www.whitehouse.gov/presidential-actions/2026/01/adjusting-imports-of-semiconductors-semiconductor-manufacturing-equipment-and-their-derivative-products-into-the-united-states/

11. Semiconductors Insight — US–China export controls: H200 whiplash; VEU status expiry Dec 2025: https://semiconductorsinsight.com/us-china-chip-export-controls-h200-2026/

12. TariffWise — Section 301 on Chinese semiconductors doubled to 50%; Section 232 phase-in — https://tariffwise.co/news/semiconductor-tariffs-2026/

13. Yole Group — Power SiC + GaN to exceed $14B in five years; automotive slowdown hit SiC in 2024–25: https://www.yolegroup.com/strategy-insights/power-sic-momentum-builds-as-the-global-supply-chain-evolves-insights-ahead-of-semicon-china-2026/

14. GlobalDeal / Finder — CHIPS direct funding (TSMC Arizona $6.6B; Intel $7.86B); EU & Japan incentives: https://globaldeal.io/finder/articles/semiconductor-incentives-post-chips-2026

15. European Commission — European Chips Act mobilised >€52B in public & private investment: https://digital-strategy.ec.europa.eu/en/policies/chips-act-2

16. JapanTalkback — Japan's cumulative ~$15B support for Rapidus 2nm programme: https://kantenna.com/topic/rapidus-6315-billion-yen-subsidy-japan-semiconductor-national-strategy

17. CNBC — Intel CEO Lip-Bu Tan on foundry as a "national treasure" and opening capacity to outside customers: https://www.cnbc.com/2026/05/18/intel-ceo-says-foundry-is-gaining-momentum-as-customer-interest-grows.html

Authors



Sonali Salaria

Director, Sales Excellence Team, LTM

Sonali Salaria is part of Sales Excellence Team managing Sales Research & Strategy. She has over 18 years of work experience in strategic consulting, advisory, and research. She has done BE in Electrical & Electronics and PGDM in Marketing.



Hinal Das

Senior Specialist, Sales Excellence Team, LTM

Hinal Das is part of Sales Excellence Team and is an experienced professional in Sales Research, Strategy & Insights across CME, Hi-Tech and RCG sectors. She has over 17 years of work experience in market research & consulting and sales strategy. She holds an MBA and a bachelor's in management studies and is based out of Mumbai.



SenthilPrabhu G

Senior Account Director, Global Sales, LTM

Senthil Prabu Govindarajan is part of the Hi-Tech Industry Business Unit, where he manages key engagements across the semiconductor portfolio. He brings over 15 years of experience in technology consulting and sales across North America and APAC. He holds a PGDM in Management from IIM Ahmedabad and is based in Phoenix, Arizona.



Siddarth Tater

Associate Director, Program & Project Management, LTM

Siddarth Tater is an experienced leader in enterprise IT transformation and Application Managed Support, with expertise in program governance, Oracle ERP/SCM, and operational analytics across the semiconductor and Hi-Tech sectors. He specializes in large-scale legacy-to-cloud migrations, capacity-aware order planning, and AMS stabilization for global clients. He translates complex operational data into leadership-ready insights and is based in Scottsdale, Arizona.

LTM — a Larsen & Toubro Group Company — is an AI-centric global technology services company and the Business Creativity partner to the world's largest enterprises.

We bring human insights and intelligent systems together to help clients create greater value at the intersection of technology and domain expertise. Our capabilities span integrated operations, transformation, and business AI — enabling new ways of working, new productivity paradigms, and new roads to value. Together with over 88,000 employees across 40 countries and our global network of partners, LTM owns outcomes for our clients, helping them not just outperform the market, but Outcreate it.

© LTM