TROY TECHNICAL

SEMICONDUCTOR BACK-END WEEKLY REPORT

September 12 - September 18, 2026 | Issue 2026-W37

Package performance, process yield, test coverage and qualified capacity

16	10
ANALYZED ARTICLES	PRIORITY THEMES

Issued 2026-09-18 | English edition | fixed approved cover artwork

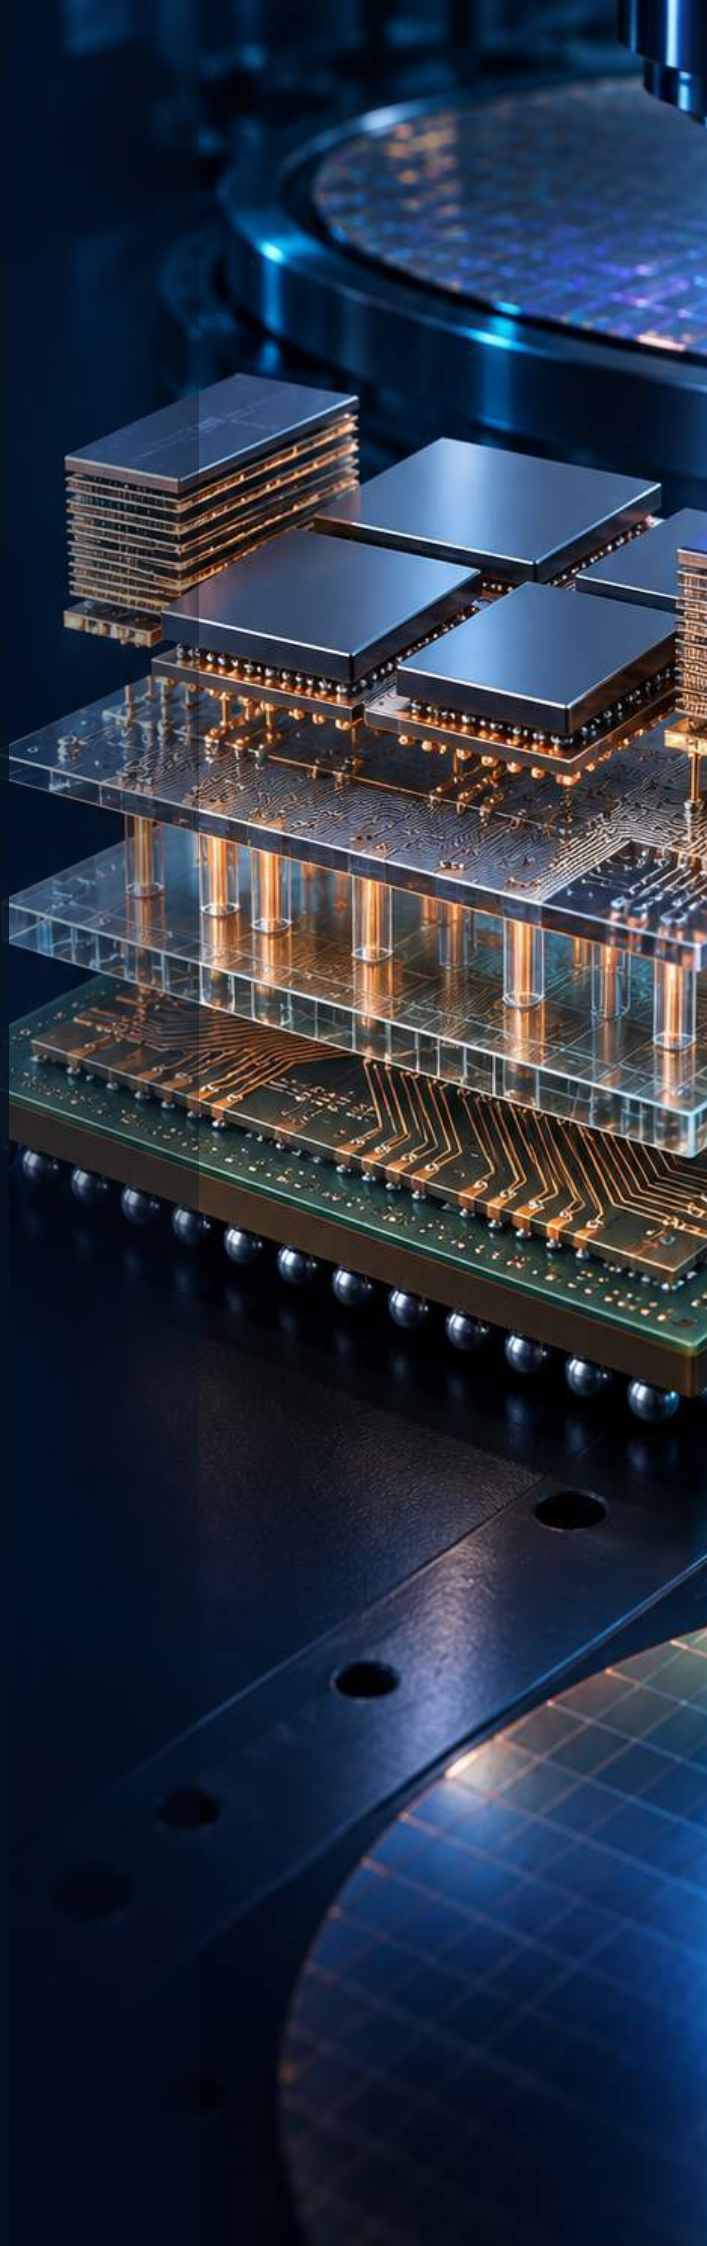

Semiconductor Back-End: the boundary moved from isolated claims to qualification evidence

16 English articles were reviewed independently from the Japanese edition. Ten themes are retained for management decisions.

SIGNAL 01 | MATERIALS

ASE's SPIL Nears Arizona Back-End Test Facility Expansion Amidst $10.5B CAPEX Surge; Amkor

ASE subsidiary SPIL is finalizing plans for a new back-end test facility in Arizona, part of ASE's $10.5 billion CAPEX for 2026, which includes $4 billion for new facilities and $6.5 billion for equipment.

SIGNAL 02 | BONDING & ASSEMBLY

AI Accelerator ABF Substrate Shortage to Exceed 40% by 2028; Ibiden, Unimicron, Samsung

The supply shortage of ABF (Ajinomoto Build-up Film) substrates for AI accelerators is projected to worsen significantly, potentially reaching approximately 10% by late 2026, 21% by 2027, and over 40% by 2028.

SIGNAL 03 | ADVANCED PACKAGE

Samsung Electronics Reorganizes Cheonan Plant for HBM Packaging, Outsourcing Mobile AP Backend Processes

Samsung Electronics is reportedly reorganizing its Cheonan plant's backend operations to prioritize High Bandwidth Memory (HBM) packaging due to surging AI demand.

SIGNAL 04 | TEST & INSPECTION

BESI Hybrid Bonding System Orders Exceed 150 Units, HBM5 Onward to Accelerate Memory

BESI has secured over 150 cumulative orders for its hybrid bonding systems, with six integrated production lines (including 30 BESI bonders) implemented by major logic customers in partnership with Applied Materials.

DECISION

Focus this week's management attention on package performance, process yield, test coverage and qualified capacity. Move only when the linked evidence can be reproduced under your own operating conditions.

THREE MANAGEMENT MOVES

1. Buyers: compare architectures using thermal, electrical, yield, repair and supply evidence.
2. Suppliers: co-develop materials, bonding, inspection and reliability windows with customers.
3. Executives: track yield, warpage, thermal cycling, test time, capacity, cost and qualification.

EVIDENCE DISCIPLINE

FACT states what the collected English source reports. BUSINESS READ and ACTION are editorial interpretations. UNKNOWN lists information that the source file does not establish. A linked article is not treated as independent confirmation of every claim.

Five numbers or evidence anchors that require conditions

$10.5 bil lion
A09

SOURCE-BOUND CONDITION
ASE subsidiary SPIL is finalizing plans for a new back-end test facility in Arizona, part of ASE's $10.5 billion CAPEX for 2026, which includes $4 billion for new facilities and $6.5 billion for equipment.

10%
A12

SOURCE-BOUND CONDITION
The supply shortage of ABF (Ajinomoto Build-up Film) substrates for AI accelerators is projected to worsen significantly, potentially reaching approximately 10% by late 2026, 21% by 2027, and over 40% by 2028.

000 w
A01

SOURCE-BOUND CONDITION
TSMC plans to double its CoWoS advanced packaging capacity from 130,000 wafers/month to 260,000 wafers/month by the end of 2028, primarily driven by surging AI chip demand.

2nm
A02

SOURCE-BOUND CONDITION
TSMC plans a significant capacity expansion for its 2nm and 3nm advanced processes to meet surging demand for AI and HPC chips.

90%
A03

SOURCE-BOUND CONDITION
According to Epoch AI estimates, just four companies-Nvidia, Google, AMD, and Amazon-consumed approximately 90% of the global HBM supply in 2025.

HOW TO USE THESE NUMBERS
Each value is shown with the condition stated in the collected English article. Do not compare values across different test methods, device sizes, operating temperatures or commercial stages without normalizing those conditions.

Where the value chain moved

Each stage combines one leading theme and one supporting theme. The report treats handoffs as evidence transfers, not decorative arrows.

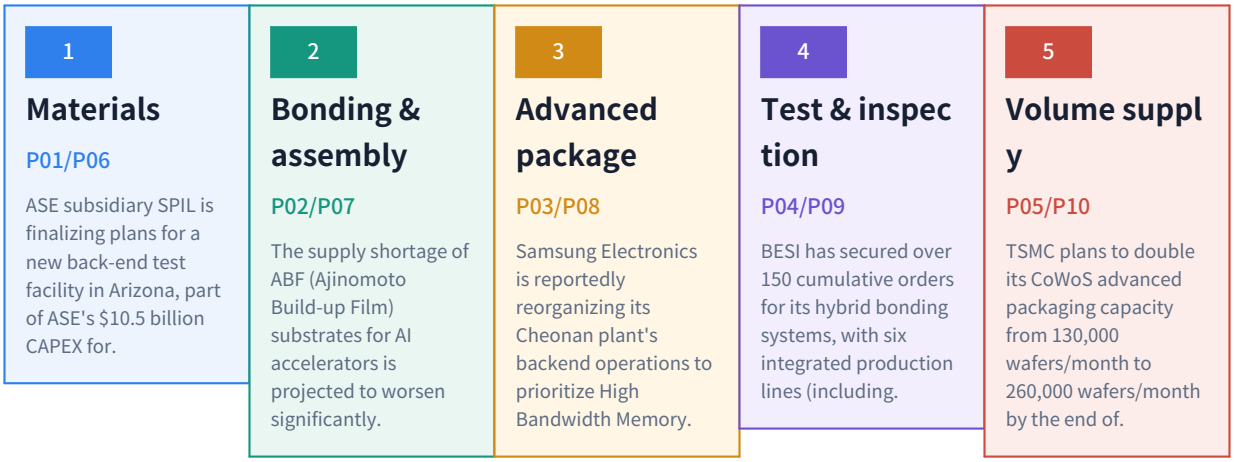

EVIDENCE REQUIRED AT EACH HANDOFF

Materials to Bonding & assembly	Transfer test conditions, acceptance criteria, failure data, owner and decision date. Recheck yield, warpage, thermal cycling, test time, capacity, cost and qualification.
Bonding & assembly to Advanced package	Transfer test conditions, acceptance criteria, failure data, owner and decision date. Recheck yield, warpage, thermal cycling, test time, capacity, cost and qualification.
Advanced package to Test & inspection	Transfer test conditions, acceptance criteria, failure data, owner and decision date. Recheck yield, warpage, thermal cycling, test time, capacity, cost and qualification.
Test & inspection to Volume supply	Transfer test conditions, acceptance criteria, failure data, owner and decision date. Recheck yield, warpage, thermal cycling, test time, capacity, cost and qualification.

COMMON CONTROL POINT

No theme moves to a commercial commitment until a named owner has reproduced the relevant evidence and documented yield, warpage, thermal cycling, test time, capacity, cost and qualification.

Priority map: maturity and business impact

P01	ASE's SPIL Nears Arizona Back-End Test Facility Expansion Amidst $10.5B CAPEX Surge; Amkor
P02	AI Accelerator ABF Substrate Shortage to Exceed 40% by 2028; Ibiden, Unimicron, Samsung
P03	Samsung Electronics Reorganizes Cheonan Plant for HBM Packaging, Outsourcing Mobile AP Backend Processes
P04	BESI Hybrid Bonding System Orders Exceed 150 Units, HBM5 Onward to Accelerate Memory
P05	TSMC to Double CoWoS Advanced Packaging Capacity to 260,000 Wafers/Month by 2028, Driving
P06	CSIS Report: AI and HBM Demand Drive New Semiconductor Shortage, Highlighting Structural Shift
P07	Nvidia Chip Sales Projected to Double Next Year, Accelerating HBM4 Transition Amid TSMC
P08	TSMC to Boost 2nm and 3nm Advanced Process Capacity by up to 22%
P09	HBM Identified as Primary Bottleneck for AI Chip Manufacturing; Nvidia, Google, AMD, Amazon
P10	AI Infrastructure Component Crunch Becomes Reality: ABF Substrate Lead Times Hit 56 Weeks

REPRODUCIBLE POSITIONING RULE

Horizontal position uses five discrete stages: Research, Prototype, Joint evaluation, Product adoption and Scale. Vertical position uses three impact levels: single use, multiple customers and multiple supply tiers. Bubble diameter reflects the editorial score inherited from the weekly selection rubric.

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P01 ASE's SPIL Nears Arizona Back-End Test Facility Expansion Amidst $10.5B CAPEX Surge; Amkor

A09 | Materials | Product adoption

Date not stated

FACT

ASE subsidiary SPIL is finalizing plans for a new back-end test facility in Arizona, part of ASE's $10.5 billion CAPEX for 2026, which includes $4 billion for new facilities and $6.5 billion for equipment. Concurrently, SPIL is expanding CoWoS advanced packaging in Douliu, Taiwan.

WHY IT MATTERS

The decision relevance is package performance, process yield, test coverage and qualified capacity. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, co-develop materials, bonding, inspection and reliability windows with customers.

ACTION

Within 30 days, compare architectures using thermal, electrical, yield, repair and supply evidence; record the evidence and owner against P01.

UNKNOWN

Still unproven or incompletely stated: yield, warpage, thermal cycling, test time, capacity, cost and qualification.

EDITORIAL SCORE 22/25 | MATURITY 4/5 | IMPACT 1/3

P02 AI Accelerator ABF Substrate Shortage to Exceed 40% by 2028; Ibiden, Unimicron, Samsung

A12 | Bonding & assembly | Scale-up

Date not stated

FACT

The supply shortage of ABF (Ajinomoto Build-up Film) substrates for AI accelerators is projected to worsen significantly, potentially reaching approximately 10% by late 2026, 21% by 2027, and over 40% by 2028. In response, Ibiden is investing ¥500 billion (approx. $3.1B) from 2026-2028 to increase production capacity 2.8-fold compared to 2024. Unimicron is dedicating a record NT$34 billion (approx.

WHY IT MATTERS

The decision relevance is package performance, process yield, test coverage and qualified capacity. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, co-develop materials, bonding, inspection and reliability windows with customers.

ACTION

Within 30 days, compare architectures using thermal, electrical, yield, repair and supply evidence; record the evidence and owner against P02.

UNKNOWN

Still unproven or incompletely stated: yield, warpage, thermal cycling, test time, capacity, cost and qualification.

EDITORIAL SCORE 22/25 | MATURITY 5/5 | IMPACT 2/3

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P03 Samsung Electronics Reorganizes Cheonan Plant for HBM Packaging, Outsourcing Mobile AP Backend Processes

A04 | Advanced package | Scale-up

Date not stated

FACT

Samsung Electronics is reportedly reorganizing its Cheonan plant's backend operations to prioritize High Bandwidth Memory (HBM) packaging due to surging AI demand. This strategic shift involves phasing out Exynos mobile application processor (AP) packaging and repurposing equipment for HBM, with the transition expected to commence next year.

WHY IT MATTERS

The decision relevance is package performance, process yield, test coverage and qualified capacity. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, co-develop materials, bonding, inspection and reliability windows with customers.

ACTION

Within 30 days, compare architectures using thermal, electrical, yield, repair and supply evidence; record the evidence and owner against P03.

UNKNOWN

Still unproven or incompletely stated: yield, warpage, thermal cycling, test time, capacity, cost and qualification.

EDITORIAL SCORE 21/25 | MATURITY 5/5 | IMPACT 2/3

P04 BESI Hybrid Bonding System Orders Exceed 150 Units, HBM5 Onward to Accelerate Memory

A11 | Test & inspection | Scale-up

Date not stated

FACT

BESI has secured over 150 cumulative orders for its hybrid bonding systems, with six integrated production lines (including 30 BESI bonders) implemented by major logic customers in partnership with Applied Materials. Currently deployed primarily for logic and Co-packaged Optics (CPO), this technology is expected to see memory adoption, specifically for HBM5 (20-layer stacking) and beyond, after HBM4E continues with existing methods, anticipating a three-year delay in memory sector penetration.

WHY IT MATTERS

The decision relevance is package performance, process yield, test coverage and qualified capacity. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, co-develop materials, bonding, inspection and reliability windows with customers.

ACTION

Within 30 days, compare architectures using thermal, electrical, yield, repair and supply evidence; record the evidence and owner against P04.

UNKNOWN

Still unproven or incompletely stated: yield, warpage, thermal cycling, test time, capacity, cost and qualification.

EDITORIAL SCORE 21/25 | MATURITY 5/5 | IMPACT 2/3

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P05 TSMC to Double CoWoS Advanced Packaging Capacity to 260,000 Wafers/Month by 2028, Driving

A01 | Volume supply | Product adoption

Date not stated

FACT

TSMC plans to double its CoWoS advanced packaging capacity from 130,000 wafers/month to 260,000 wafers/month by the end of 2028, primarily driven by surging AI chip demand. This expansion, centered on its AP7 facility in Taiwan and the Arizona plant, aims to alleviate bottlenecks.

WHY IT MATTERS

The decision relevance is package performance, process yield, test coverage and qualified capacity. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, co-develop materials, bonding, inspection and reliability windows with customers.

ACTION

Within 30 days, compare architectures using thermal, electrical, yield, repair and supply evidence; record the evidence and owner against P05.

UNKNOWN

Still unproven or incompletely stated: yield, warpage, thermal cycling, test time, capacity, cost and qualification.

EDITORIAL SCORE 21/25 | MATURITY 4/5 | IMPACT 2/3

P06 CSIS Report: AI and HBM Demand Drive New Semiconductor Shortage, Highlighting Structural Shift

A08 | Materials | Product adoption

Date not stated

FACT

A CSIS report indicates that the current memory shortage reflects both structural demand growth from AI and cyclical market fluctuations. The rapid expansion of AI data centers has led to a surge in demand for HBM (High Bandwidth Memory), essential for AI workloads, which in turn strains conventional DRAM supply due to HBM's need for advanced packaging. The concentration of commercial-scale HBM production in Asia raises supply chain concentration risks and national security concerns, even with increasing U.S.

WHY IT MATTERS

The decision relevance is package performance, process yield, test coverage and qualified capacity. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, co-develop materials, bonding, inspection and reliability windows with customers.

ACTION

Within 30 days, compare architectures using thermal, electrical, yield, repair and supply evidence; record the evidence and owner against P06.

UNKNOWN

Still unproven or incompletely stated: yield, warpage, thermal cycling, test time, capacity, cost and qualification.

EDITORIAL SCORE 21/25 | MATURITY 4/5 | IMPACT 3/3

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P07 Nvidia Chip Sales Projected to Double Next Year, Accelerating HBM4 Transition Amid TSMC

A06 | Bonding & assembly | Product adoption

Date not stated

FACT

Nvidia CEO Jensen Huang forecasts a doubling of Nvidia chip sales next year and an accelerated transition to HBM4, with TSMC's CoWoS advanced packaging identified as a primary shipping bottleneck. CoWoS is crucial for combining Nvidia GPUs and HBM into single packages. TSMC plans to double its monthly CoWoS capacity from approximately 130,000 wafers by the end of 2026 to 260,000 wafers by the end of 2028, an expansion critical for future supply.

WHY IT MATTERS

The decision relevance is package performance, process yield, test coverage and qualified capacity. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, co-develop materials, bonding, inspection and reliability windows with customers.

ACTION

Within 30 days, compare architectures using thermal, electrical, yield, repair and supply evidence; record the evidence and owner against P07.

UNKNOWN

Still unproven or incompletely stated: yield, warpage, thermal cycling, test time, capacity, cost and qualification.

EDITORIAL SCORE 20/25 | MATURITY 4/5 | IMPACT 2/3

P08 TSMC to Boost 2nm and 3nm Advanced Process Capacity by up to 22%

A02 | Advanced package | Product adoption

Date not stated

FACT

TSMC plans a significant capacity expansion for its 2nm and 3nm advanced processes to meet surging demand for AI and HPC chips. The 2nm monthly capacity is projected to increase by approximately 22% to 110,000 wafers by mid-2027 from 90,000 wafers at the end of 2026, while 3nm capacity will rise by over 16% to 210,000 wafers from 180,000. This expansion, coupled with a planned doubling of CoWoS advanced packaging capacity by 2028, aims to bolster the comprehensive supply for AI chips.

WHY IT MATTERS

The decision relevance is package performance, process yield, test coverage and qualified capacity. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, co-develop materials, bonding, inspection and reliability windows with customers.

ACTION

Within 30 days, compare architectures using thermal, electrical, yield, repair and supply evidence; record the evidence and owner against P08.

UNKNOWN

Still unproven or incompletely stated: yield, warpage, thermal cycling, test time, capacity, cost and qualification.

EDITORIAL SCORE 20/25 | MATURITY 4/5 | IMPACT 2/3

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P09 HBM Identified as Primary Bottleneck for AI Chip Manufacturing; Nvidia, Google, AMD, Amazon

A03 | Test & inspection | Product adoption

Date not stated

FACT

High Bandwidth Memory (HBM) has emerged as the most significant bottleneck in AI chip manufacturing, surpassing advanced packaging technologies like CoWoS. According to Epoch AI estimates, just four companies-Nvidia, Google, AMD, and Amazon-consumed approximately 90% of the global HBM supply in 2025. While CoWoS capacity has seen some relief through TSMC's expansions, HBM consumption nearly matches the entire global supply, making memory scarcity the most critical manufacturing constraint.

WHY IT MATTERS

The decision relevance is package performance, process yield, test coverage and qualified capacity. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, co-develop materials, bonding, inspection and reliability windows with customers.

ACTION

Within 30 days, compare architectures using thermal, electrical, yield, repair and supply evidence; record the evidence and owner against P09.

UNKNOWN

Still unproven or incompletely stated: yield, warpage, thermal cycling, test time, capacity, cost and qualification.

EDITORIAL SCORE 20/25 | MATURITY 4/5 | IMPACT 3/3

P10 AI Infrastructure Component Crunch Becomes Reality: ABF Substrate Lead Times Hit 56 Weeks

A14 | Volume supply | Product adoption

Date not stated

FACT

The supply shortage for AI infrastructure components is becoming entrenched; TrendForce reports ABF substrate lead times surged to 48-56 weeks, approximately 4 to 4.7 times the equilibrium. HDDs reached 50 weeks, DRAM 20 weeks, and MLCCs 30 weeks, with all five tracked categories (excluding GPUs) in shortage. Semiconductor equipment parts are particularly critical, with some Japanese-made crucial components facing lead times up to 40 months.

WHY IT MATTERS

The decision relevance is package performance, process yield, test coverage and qualified capacity. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, co-develop materials, bonding, inspection and reliability windows with customers.

ACTION

Within 30 days, compare architectures using thermal, electrical, yield, repair and supply evidence; record the evidence and owner against P10.

UNKNOWN

Still unproven or incompletely stated: yield, warpage, thermal cycling, test time, capacity, cost and qualification.

EDITORIAL SCORE 20/25 | MATURITY 4/5 | IMPACT 2/3

Playbook and 0-5 year roadmap

STAKEHOLDER	NOW | 0-30 DAYS	NEXT | 1-12 MONTHS	LATER | 1-5 YEARS
Operating companies	compare architectures using thermal, electrical, yield, repair and supply evidence	Run production-like qualification with explicit acceptance and stop criteria.	Build a portfolio and architecture that can absorb technology and supplier changes.
Materials, equipment and technology suppliers	co-develop materials, bonding, inspection and reliability windows with customers	Create shared reliability, process-window and failure-analysis data with customers.	Standardize interfaces, traceability, lifecycle support and recovery services.
Trading, investment and legal teams	Map yield, warpage, thermal cycling, test time, capacity, cost and qualification.	Contract for capacity, quality, change notice, liability, alternatives and exit.	Diversify regional, technical and customer concentration risk.

ROADMAP

NOW | 0-30 DAYS
Evidence ledger
Owner: Strategy + quality
Deliverable: claim, source, condition, owner

NOW | 0-30 DAYS
Risk screen
Owner: Procurement + legal
Deliverable: unknowns, alternatives, stop

NEXT | 1-12 MONTHS
Joint qualification
Owner: R&D + supplier
Deliverable: process window and failure data

NEXT | 1-12 MONTHS
Commercial gate
Owner: Business + finance
Deliverable: unit economics and capacity

LATER | 1-5 YEARS
Platform strategy
Owner: Executive team
Deliverable: portfolio and interface roadmap

LATER | 1-5 YEARS
Service layer
Owner: Operations + channel
Deliverable: monitoring, maintenance and recovery

GO / NO-GO GATES

GATE	OWNER	REQUIRED EVIDENCE	NO-GO CONDITION
G1 Evidence	**Quality**	Source, test conditions and reproducibility	No traceable evidence
G2 Integration	**R&D**	Interface, process window and failure analysis	Cannot meet system conditions
G3 Commercial	**Business**	Capacity, total cost, contract and alternatives	Economics or supply cannot be supported
G4 Lifecycle	**Operations**	Monitoring, maintenance, change notice and exit	No accountable operating plan

ANALYZED ARTICLES | ITEMS 01-08

Every title and URL is displayed in full. URLs wrap without shortening and remain clickable. A verified-reference label means the translated HTML did not retain its original source URL.

ID	FULL ARTICLE TITLE	FULL SOURCE / VERIFIED REFERENCE URL - CLICKABLE
A01	TSMC to Double CoWoS Advanced Packaging Capacity to 260,000 Wafers/Month by 2028, Driving OSAT and Intel Orders Amid Supply Crunch	https://finance.biggo.com/news/adcec901-53b3-453f-bf00-8373ebb1bb97
A02	TSMC to Boost 2nm and 3nm Advanced Process Capacity by up to 22% by Mid-2027 to Meet Soaring AI/HPC Demand	https://finance.biggo.com/news/b8e1ef85-db81-4a2a-a33b-b419fe0c005f
A03	HBM Identified as Primary Bottleneck for AI Chip Manufacturing; Nvidia, Google, AMD, Amazon Consumed 90% of Global Supply in 2025	https://newmarketpitch.com/blogs/news/ai-chip-challenges
A04	Samsung Electronics Reorganizes Cheonan Plant for HBM Packaging, Outsourcing Mobile AP Backend Processes to Boost AI Memory Capacity	https://www.trendforce.com/news/2026/09/16/news-samsung-reportedly-shifts-cheonan-capacity-to-hbm-packaging-outsource-conventional-memory-modules/
A05	FCS Group Secures Exclusive Taiwan/Southeast Asia Distribution Rights for CKST Glass Substrate Equipment, Entering AI Advanced Packaging Market	https://finance.biggo.com/news/32947af3-c665-4534-8c4b-eb2f255278f7
A06	Nvidia Chip Sales Projected to Double Next Year, Accelerating HBM4 Transition Amid TSMC CoWoS Packaging Bottleneck	https://finance.biggo.com/news/d81ffb54-2ad9-4684-805d-14a08d4f7eeb
A07	Google Reportedly Considering CoWoS as Backup for 2027 TPUs Amid Intel EMIB-T Substrate Yield Lag Reports	https://www.digitimes.com/news/a20260917PD208/google-tpu-cowos-substrate-2027.html
A08	CSIS Report: AI and HBM Demand Drive New Semiconductor Shortage, Highlighting Structural Shift and Supply Chain Security Concerns	https://www.csis.org/analysis/beyond-memory-cycle-ai-hbm-and-new-semiconductor-shortage

ANALYZED ARTICLES | ITEMS 09-16

Every title and URL is displayed in full. URLs wrap without shortening and remain clickable. A verified-reference label means the translated HTML did not retain its original source URL.

ID	FULL ARTICLE TITLE	FULL SOURCE / VERIFIED REFERENCE URL - CLICKABLE
A09	ASE's SPIL Nears Arizona Back-End Test Facility Expansion Amidst $10.5B CAPEX Surge; Amkor Boosts US Packaging Investment to $12B	https://www.trendforce.com/news/2026/09/15/news-ases-spil-reportedly-nears-arizona-expansion-as-tsmc-amkor-step-up-u-s-advanced-packaging-push/
A10	Samsung's zHBM Concept Promises 70% Lower I/O Power, 2.3x Bandwidth vs. HBM5 via Hybrid Bonding for Edge AI	https://www.tomshardware.com/tech-industry/semiconductors/piecemakers-bets-edge-ai-devices-will-diverge-from-reliance-on-hbm-custom-designed-memory-fuses-dram-stack-directly-to-the-processor-using-hybrid-bonding
A11	BESI Hybrid Bonding System Orders Exceed 150 Units, HBM5 Onward to Accelerate Memory Adoption	https://photoncap.net/p/investment-map-22-listed-companies
A12	AI Accelerator ABF Substrate Shortage to Exceed 40% by 2028; Ibiden, Unimicron, Samsung Electro-Mechanics Rush Billions in Capacity Expansion	https://www.tomshardware.com/tech-industry/semiconductors/the-state-of-abf-substrates-in-data-center-silicon-in-2026-solving-the-supply-crunch-and-material-wall-beneath-every-ai-accelerator
A13	Samsung Electro-Mechanics Unveils Next-Gen Package Substrate Technologies for AI & Servers at KPCA Show 2026, Highlighting Warp Reduction & High-Density Interconnects	https://www.semiconductor-digest.com/samsung-electro-mechanics-showcases-next-generation-semiconductor-package-substrate-technologies-for-ai-and-servers-at-kpca-show-2026/
A14	AI Infrastructure Component Crunch Becomes Reality: ABF Substrate Lead Times Hit 56 Weeks, Equipment Parts Up to 40 Months	https://finance.biggo.com/news/02cf5c8f-c58d-45f8-a7a3-8fe37716bc8f
A15	Glass Substrates & TGV Technology Set to Reshape 2027 Advanced Packaging Race; US, Japan, Korea, Taiwan, China Accelerate Development	https://www.digitimes.com/news/a20260916PD221/ai-chip-2027-substrate-technology-materials.html
A16	Samsung Electro-Mechanics and Qualcomm Jointly Develop Organic Bridge for Advanced 2.1D AI Semiconductor Packaging	https://www.thelec.net/news/articleView.html?idxno=13907

Semiconductor_BackEnd — Selected Articles

Date: 2026-09-20

Articles: 16

Table of Contents

#01 TSMC to Double CoWoS Advanced Packaging Capacity to 260,000 Wafers/Month by 2028, Driving OSAT and Intel Orders Amid Supply Crunch

Published September 13, 2026 BigGo Finance Taiwan

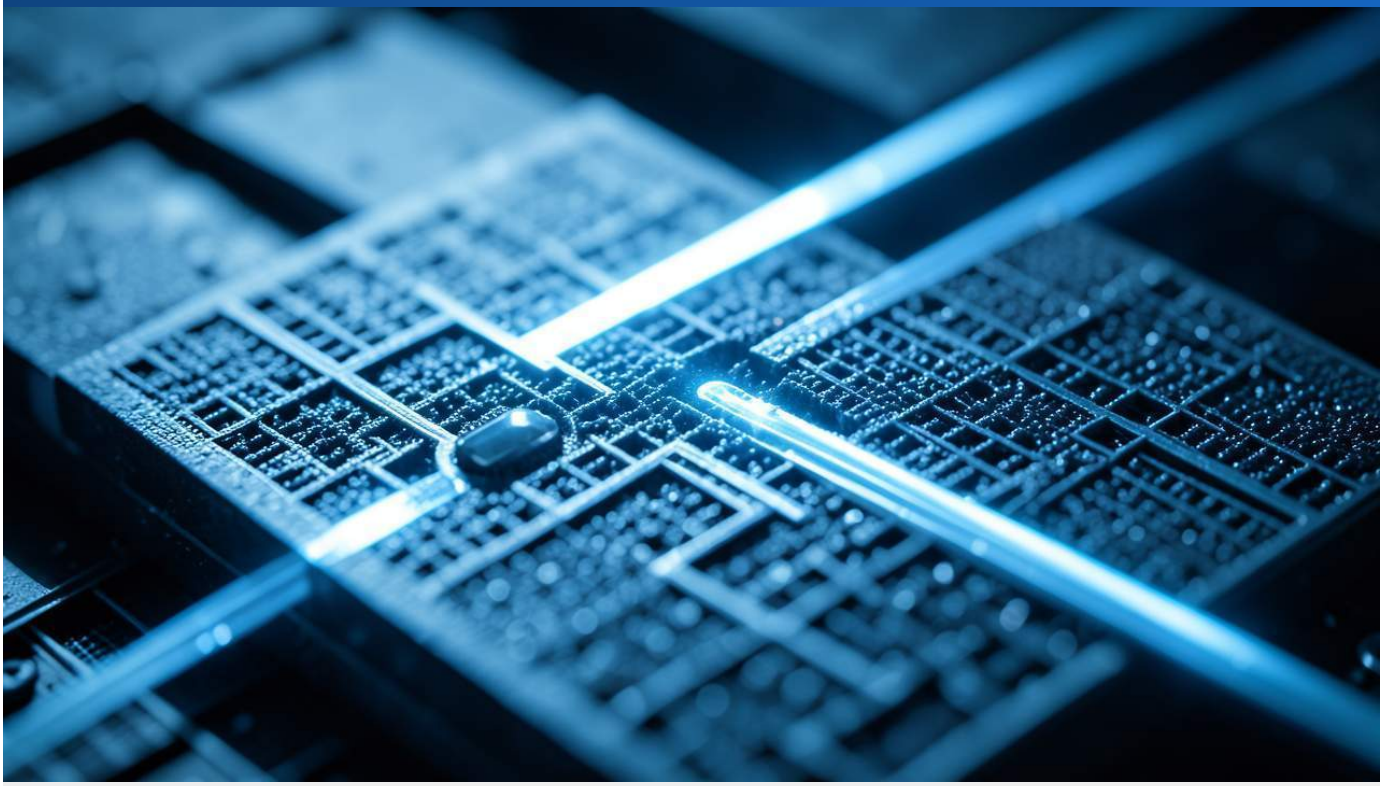

OVERVIEW

TSMC plans to double its CoWoS advanced packaging capacity from 130,000 wafers/month to 260,000 wafers/month by the end of 2028, primarily driven by surging AI chip demand. This expansion, centered on its AP7 facility in Taiwan and the Arizona plant, aims to alleviate bottlenecks. The persistent demand exceeding TSMC's supply has led to spillover orders for OSAT providers like ASE and Amkor, while Intel positions its EMIB-T technology to capture custom AI chip packaging, targeting 40,000-45,000 wafer-equivalent capacity by 2028.

IN DEPTH

Key Findings

TSMC, the world's leading contract chipmaker, has announced ambitious plans to double its CoWoS (Chip-on-Wafer-on-Substrate) advanced packaging capacity to 260,000 wafers per month by the end of 2028. This significant expansion, up from an estimated 130,000 wafers per month at the end of 2026, is a direct response to the unprecedented demand for AI chips and high-performance computing (HPC).

Technical / Clinical Details

CoWoS is a crucial 3D stacking technology that integrates multiple chips, such as GPUs and High Bandwidth Memory (HBM), onto a single package, enabling the high-density and high-performance requirements of modern AI accelerators. The expansion will primarily leverage TSMC's AP7 facility in Taiwan and its upcoming Arizona plant in the United States. Due to TSMC's existing supply constraints, the robust AI chip demand is funneling orders to Outsourced Semiconductor Assembly and Test (OSAT) providers including ASE Technology Holding, Amkor Technology, United Microelectronics Corp (UMC), and Vanguard International Semiconductor Corp (VISC), who are also expanding their advanced packaging capabilities. Intel is aggressively entering the custom AI chip packaging market with its EMIB-T (Embedded Multi-die Interconnect Bridge - Tiled) technology, projected to reach a capacity equivalent to 40,000-45,000 wafers per month by 2028. This positions Intel as a key player for custom AI accelerators, particularly for hyperscalers like Google and Amazon, with MediaTek also reportedly praising EMIB packaging as "very good."

Background & Context

The explosive growth of the AI market has created a structural shortage in the semiconductor supply chain, with advanced packaging — specifically CoWoS and HBM — identified as a primary bottleneck. The current CoWoS supply-demand gap is estimated at around 20%, with projections suggesting it could exceed 30% by 2027. NVIDIA CEO Jensen Huang has consistently highlighted advanced packaging, HBM, and interconnects as critical supply chain constraints. TSMC's massive investment reflects the industry's collective effort to ramp up capacity for these indispensable components, which are foundational to the ongoing AI infrastructure buildout. This dynamic also underscores the strategic importance of HBM, with estimates suggesting four companies (Nvidia, Google, AMD, Amazon) consumed approximately 90% of the global HBM supply in 2025.

Strategic Significance & Outlook

The strategic capacity ramp-up by TSMC, complemented by OSAT providers and Intel, is vital for sustaining the rapid advancement and deployment of AI technologies. This expansion not only addresses immediate supply shortfalls but also paves the way for future innovations in chip design and integration. As the industry moves towards next-generation packaging solutions, such as glass substrates, the ability to scale high-volume advanced packaging will differentiate market leaders. The alleviation of these bottlenecks is expected to accelerate the introduction of more sophisticated AI chips, further fueling growth across the semiconductor ecosystem and AI-driven industries globally.

Source: https://finance.biggo.com/news/adcec901-53b3-453f-bf00-8373ebb1bb97

Collected: September 18, 2026 | Automated Research System (Gemini API)

#02 TSMC to Boost 2nm and 3nm Advanced Process Capacity by up to 22% by Mid-2027 to Meet Soaring AI/HPC Demand

Published September 15, 2026 BigGo Finance Taiwan

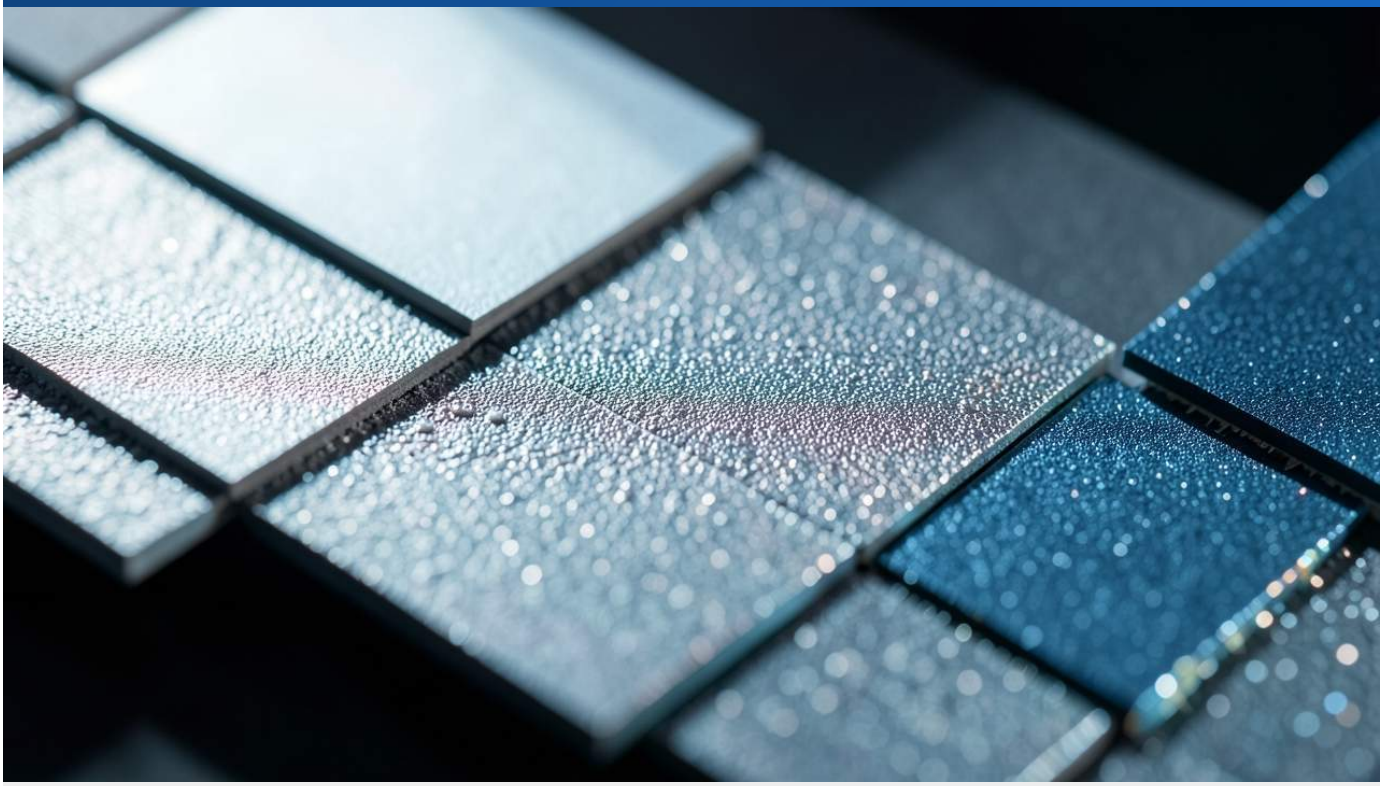

OVERVIEW

TSMC plans a significant capacity expansion for its 2nm and 3nm advanced processes to meet surging demand for AI and HPC chips. The 2nm monthly capacity is projected to increase by approximately 22% to 110,000 wafers by mid-2027 from 90,000 wafers at the end of 2026, while 3nm capacity will rise by over 16% to 210,000 wafers from 180,000. This expansion, coupled with a planned doubling of CoWoS advanced packaging capacity by 2028, aims to bolster the comprehensive supply for AI chips.

IN DEPTH

Key Findings

TSMC, a pivotal player in the global semiconductor industry, is set to significantly ramp up its production capacity for leading-edge 2nm and 3nm process technologies by mid-2027. This aggressive expansion, with increases of up to 22% for 2nm and over 16% for 3nm, directly addresses the escalating global demand for AI and High-Performance Computing (HPC) chips, which are central to the next wave of technological innovation.

Technical / Clinical Details

Specifically, TSMC's monthly production capacity for 2nm process technology is projected to grow from 90,000 wafers at the end of 2026 to approximately 110,000 wafers by mid-2027, representing a 22% increase. Concurrently, 3nm process capacity will expand from 180,000 wafers to 210,000 wafers, an increase of over 16%. These advanced nodes are critical for manufacturing the most powerful and energy-efficient AI accelerators, CPUs, and GPUs. This fabrication capacity boost is synergistically aligned with TSMC's plans to double its CoWoS (Chip-on-Wafer-on-Substrate) advanced packaging capacity to 260,000 wafers per month by the end of 2028, effectively tackling critical bottlenecks in advanced chip integration. Meanwhile, Intel's EMIB-T (Embedded Multi-die Interconnect Bridge - Tiled) packaging technology is also gaining traction, reportedly well-suited for custom AI accelerators from major clients like Google and Amazon, with MediaTek also expressing positive feedback on EMIB. Intel aims for 40,000-45,000 wafer-equivalent EMIB-T capacity by 2028.

Background & Context

The burgeoning complexity of AI models and the exponential growth of data require ever more sophisticated and power-efficient processing units. This surge has placed immense pressure on the supply chain for leading-edge silicon fabrication and advanced packaging. As the preeminent foundry, TSMC's strategic investments are crucial for meeting this demand and supporting the product roadmaps of tech giants like NVIDIA, Apple, and Qualcomm. The current landscape highlights HBM (High Bandwidth Memory) and CoWoS packaging as key constraints in the AI chip supply chain, making TSMC's holistic approach to expanding both front-end (wafer fabrication) and back-end (packaging) capacities strategically vital for stabilizing the ecosystem.

Strategic Significance & Outlook

TSMC's simultaneous expansion of both advanced process nodes and CoWoS packaging capacity is poised to have a profound impact across the entire AI chip ecosystem. By stabilizing and increasing supply, more companies will be enabled to develop and deploy groundbreaking AI products, broadening the application of AI technologies across various industries. The competitive landscape, further shaped by Intel's EMIB-T advancements and the contributions of OSAT providers, indicates an unprecedented era of investment and innovation in semiconductors. This intensified competition is set to define the future of computing, pushing the boundaries of what AI can achieve.

Source: https://finance.biggo.com/news/b8e1ef85-db81-4a2a-a33b-b419fe0c005f

Collected: September 18, 2026 | Automated Research System (Gemini API)

#03 HBM Identified as Primary Bottleneck for AI Chip Manufacturing; Nvidia, Google, AMD, Amazon Consumed 90% of Global Supply in 2025

Published September 10, 2026 New Market Pitch USA

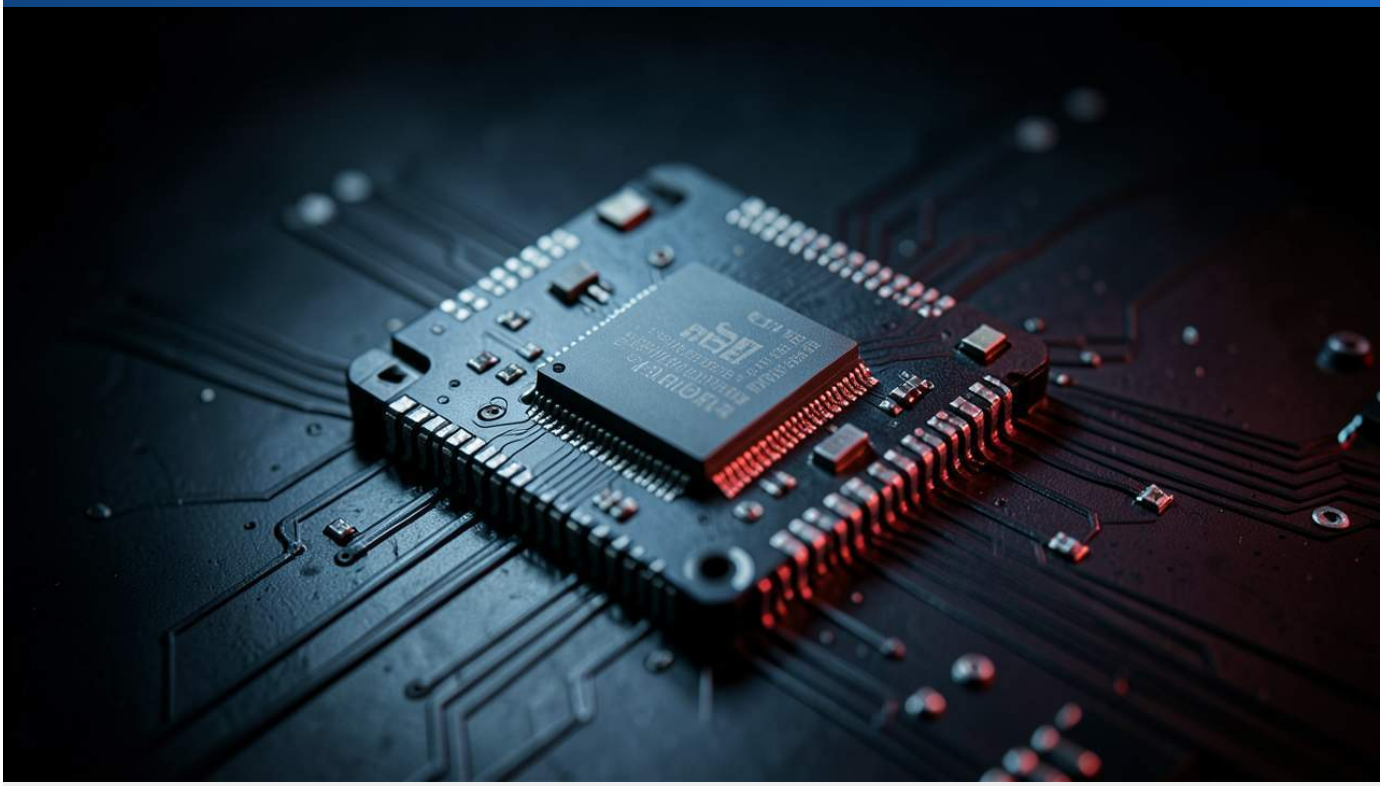

OVERVIEW

High Bandwidth Memory (HBM) has emerged as the most significant bottleneck in AI chip manufacturing, surpassing advanced packaging technologies like CoWoS. According to Epoch AI estimates, just four companies—Nvidia, Google, AMD, and Amazon—consumed approximately 90% of the global HBM supply in 2025. While CoWoS capacity has seen some relief through TSMC's expansions, HBM consumption nearly matches the entire global supply, making memory scarcity the most critical manufacturing constraint.

IN DEPTH

Key Findings

The most critical bottleneck in the manufacturing of AI chips has shifted from advanced packaging, such as CoWoS, to High Bandwidth Memory (HBM). Epoch AI estimates that in 2025, a mere four companies—Nvidia, Google, AMD, and Amazon—collectively consumed approximately 90% of the world's HBM supply.

Technical / Clinical Details

HBM is an indispensable high-speed, high-density memory technology crucial for AI accelerators and high-performance GPUs. Its stacked architecture delivers significantly greater bandwidth compared to conventional DRAM, which is vital for the intensive data processing requirements of AI workloads. While advanced packaging like CoWoS remains tight, extensive capacity expansions by leading foundries such as TSMC have provided some relief. However, HBM consumption has virtually reached the entirety of the global supply, indicating a structural challenge that demands not only manufacturing capacity expansion but also broader supply chain adjustments and technological innovation. The fundamental challenge lies in scaling HBM production to keep pace with the exponential growth in demand from leading AI developers.

Background & Context

The rapid expansion of AI infrastructure has created unprecedented demand across the semiconductor industry, leading to a succession of identified bottlenecks. Previously, leading-edge process nodes and CoWoS packaging were primary concerns, but now, the scarcity of HBM has become the most critical constraint. AI leaders like NVIDIA are aggressively expanding their data centers, and HBM is key to maximizing GPU performance within these facilities. The overwhelming dominance of HBM consumption by a few large corporations poses significant challenges for new entrants and smaller firms in the AI chip development space, potentially concentrating market power.

Strategic Significance & Outlook

The HBM supply crunch will directly impact the pace of AI chip development and deployment, making HBM manufacturing capacity expansion and efficiency improvements a top priority for the coming years. Memory manufacturers are expected to accelerate investments in HBM production and facilitate the transition to next-generation HBM (e.g., HBM4). Additionally, AI chip designers may explore novel architectures that optimize HBM utilization and investigate alternative memory solutions. Diversifying HBM supply sources and geographically dispersing production facilities will also be crucial for enhancing supply chain resilience and mitigating future bottlenecks, ensuring sustainable growth for the AI industry.

Source: https://newmarketpitch.com/blogs/news/ai-chip-challenges

#04 Samsung Electronics Reorganizes Cheonan Plant for HBM Packaging, Outsourcing Mobile AP Backend Processes to Boost AI Memory Capacity

Published September 16, 2026 TrendForce South Korea

OVERVIEW

Samsung Electronics is reportedly reorganizing its Cheonan plant's backend operations to prioritize High Bandwidth Memory (HBM) packaging due to surging AI demand. This strategic shift involves phasing out Exynos mobile application processor (AP) packaging and repurposing equipment for HBM, with the transition expected to commence next year. Samsung's 10 trillion won investment in Cheonan's HBM facilities and early mass production of HBM4 contributed to a significant increase in DRAM bit shipments in Q2 2026, marking a critical move to secure leadership in the AI memory market.

IN DEPTH

Key Findings

Samsung Electronics is reportedly reconfiguring its backend operations at the Cheonan plant to significantly shift packaging capacity towards High Bandwidth Memory (HBM), driven by the escalating demand for AI chips. This strategic move aims to solidify Samsung's leadership in the burgeoning AI memory market.

Technical / Clinical Details

According to reports, Samsung plans to gradually scale back the packaging of Exynos mobile application processors (APs) and reallocate the associated equipment for HBM packaging. This transition is anticipated to begin in earnest next year. HBM is a crucial advanced memory technology for AI accelerators and high-performance GPUs, offering significantly higher bandwidth and lower power consumption than conventional DRAM through its stacked architecture. Samsung has already invested 10 trillion won (approximately $7.5 billion USD) in its HBM facilities in Cheonan, and the early mass production and shipment of HBM4 (4th-generation HBM) contributed to a substantial increase in DRAM bit shipments in Q2 2026. This reorganization signifies a strategic pivot from commoditized products like mobile APs to high-value, AI-specific memory solutions, reflecting a broader industry trend towards specialization in critical AI components.

Background & Context

The increasing complexity and scale of AI models have led to an exponential demand for specialized memory technologies like HBM. Major AI chip manufacturers, such as NVIDIA, recognize that securing sufficient HBM supply is one of the primary bottlenecks determining the production volume of AI accelerators. Samsung's move is a direct response to this market shift, aiming to expand its share in the HBM market against key competitors like SK Hynix and Micron. As demand for mobile APs moderates and HBM demand skyrockets, reallocating resources becomes an imperative strategic decision for corporate profitability and future growth.

Strategic Significance & Outlook

Samsung's large-scale shift towards HBM packaging will have a significant impact on the entire AI chip supply chain. Stabilizing HBM supply is crucial for boosting the production capacity of AI accelerators, which in turn could accelerate the widespread adoption of AI technologies. Moving forward, Samsung is expected to pursue further innovations in HBM technology, aiming to maintain its competitive edge in HBM4 and subsequent generations. This investment and strategic reorientation represent a vital step for Samsung to preserve its position as a memory market leader and establish new growth engines in the era of AI.

Source: https://www.trendforce.com/news/2026/09/16/news-samsung-reportedly-shifts-cheonan-capacity-to-hbm-packaging-outsource-conventional-memory-modules/

Collected: September 18, 2026 | Automated Research System (Gemini API)

#05 FCS Group Secures Exclusive Taiwan/Southeast Asia Distribution Rights for CKST Glass Substrate Equipment, Entering AI Advanced Packaging Market

Published September 14, 2026 BigGo Finance Taiwan

OVERVIEW

FCS Group, an injection molding machine manufacturer, has secured exclusive distribution rights for CKST (China Korea Smart Technology) glass substrate advanced packaging equipment in Taiwan and Southeast Asia. This partnership marks FCS Group's formal entry into the next-generation packaging market driven by AI chips and HPC demand. The portfolio includes VP, VCP, and SPP electroplating equipment, with unit prices ranging from $4-6 million and gross profit margins exceeding 50%, signaling an acceleration in the adoption of glass substrate technology for advanced packaging.

IN DEPTH

Key Findings

FCS Group, an established injection molding machine manufacturer, has officially entered the advanced packaging market for AI chips and High-Performance Computing (HPC) by securing exclusive distribution rights in Taiwan and Southeast Asia for glass substrate advanced packaging equipment from CKST (China Korea Smart Technology), a South Korean-Chinese joint venture. This strategic alliance marks a significant diversification and expansion for FCS Group into a high-growth sector.

Technical / Clinical Details

The product portfolio for which FCS Group now holds exclusive distribution includes a range of electroplating equipment such as VP (Vertical Plating), VCP (Vertical Continuous Plating), and SPP (Semi-auto Plating Process) systems. These machines are crucial for the manufacturing processes involving glass substrates in advanced packaging. Glass substrates are gaining prominence as a next-generation material due to their superior electrical properties, thermal stability, and dimensional precision compared to traditional organic substrates. They enable finer line-width patterns and higher-density interconnects, which are essential for the increasingly miniaturized and stacked AI and HPC chip packages. Each unit of this equipment commands a high price point, ranging from $4 million to $6 million, and is expected to yield gross profit margins exceeding 50%. This high profitability underscores the robust demand for cutting-edge semiconductor manufacturing equipment and, specifically, the strong market interest in glass substrate technology.

Background & Context

The rapid evolution of AI chips and HPC has created both challenges and unprecedented opportunities in semiconductor packaging. Conventional packaging materials are encountering limitations with increased chip integration, leading to issues such as signal delay, power loss, and thermal management. Consequently, there is a pressing need for higher-performance substrate materials. Glass substrates are considered highly promising for achieving ultra-fine wiring and high-density interconnections, with major semiconductor players like Intel actively exploring their adoption. The entry of an injection molding machine manufacturer like FCS Group into the semiconductor backend equipment sector exemplifies a broader industry trend of cross-domain technological convergence and diversification of the supply chain, as companies seek to capitalize on new material and process innovations.

Strategic Significance & Outlook

The partnership between FCS Group and CKST is expected to accelerate the adoption and proliferation of glass substrate advanced packaging technology across Taiwan and Southeast Asia. As demand for AI chips and HPC continues its upward trajectory, the integration of glass substrates has the potential to establish new benchmarks for packaging performance. This development will further stimulate innovation in materials and equipment for semiconductor backend processes, ultimately contributing to enhanced performance and cost efficiency for AI chips. For FCS Group, this represents a strategically critical step to enter a high-growth market and diversify its revenue streams, positioning it for long-term success in the evolving semiconductor landscape.

Source: https://finance.biggo.com/news/32947af3-c665-4534-8c4b-eb2f255278f7

Collected: September 18, 2026 | Automated Research System (Gemini API)

#06 Nvidia Chip Sales Projected to Double Next Year, Accelerating HBM4 Transition Amid TSMC CoWoS Packaging Bottleneck

Published September 18, 2026　(情報元の名称なし)　USA

OVERVIEW

Nvidia CEO Jensen Huang forecasts a doubling of Nvidia chip sales next year and an accelerated transition to HBM4, with TSMC's CoWoS advanced packaging identified as a primary shipping bottleneck. CoWoS is crucial for combining Nvidia GPUs and HBM into single packages. TSMC plans to double its monthly CoWoS capacity from approximately 130,000 wafers by the end of 2026 to 260,000 wafers by the end of 2028, an expansion critical for future supply. This announcement highlights the ongoing momentum of AI infrastructure investments.

IN DEPTH

Key Findings

Nvidia CEO Jensen Huang has projected that Nvidia's chip sales will double next year, concurrently accelerating the transition to HBM4 (4th-generation High Bandwidth Memory), a critical component for high-performance computing and AI applications. However, a primary constraint for this rapid growth remains TSMC's CoWoS (Chip-on-Wafer-on-Substrate) advanced packaging capacity, identified as the main shipping bottleneck.

Technical / Clinical Details

Nvidia's graphics processing units (GPUs) are manufactured at TSMC, utilizing the CoWoS process to integrate the GPU with HBM in a single, unified package. This CoWoS technology is indispensable for achieving the high integration density and superior data bandwidth required to maximize the performance of AI accelerators. Currently, CoWoS production capacity lags behind the insatiable demand for Nvidia's products, creating a significant shipping bottleneck. In response, TSMC is actively executing plans to substantially expand its monthly CoWoS capacity, aiming to double it from approximately 130,000 wafers by the end of 2026 to around 260,000 wafers by the end of 2028. This capacity augmentation is paramount for Nvidia to meet its future shipment targets and ensure a smooth transition to HBM4, enabling more powerful and efficient AI systems.

Background & Context

The explosive growth of AI has driven unprecedented demand for high-performance GPUs and their associated HBM. As a leader in the AI chip market, Nvidia is a primary beneficiary of this demand. CEO Huang's prediction of doubled chip sales underscores that global investment in AI infrastructure is still in its early phases, with high levels of spending expected to continue for the long term. The CoWoS bottleneck in the supply chain affects not only Nvidia but also other AI chip developers, making large-scale investments in packaging capacity by TSMC and other OSAT (Outsourced Semiconductor Assembly and Test) providers an urgent industry-wide priority.

Strategic Significance & Outlook

Nvidia's robust growth forecast and accelerated HBM4 transition will have a profound impact across the semiconductor industry, particularly in advanced packaging and memory sectors. If TSMC's plan to double CoWoS capacity proceeds as scheduled, Nvidia's shipping bottlenecks are expected to ease progressively, allowing more AI chips to reach the market. This will further accelerate the proliferation of AI technologies, fostering innovation in diverse fields such as autonomous driving, robotics, and advanced data analytics. The key challenge ahead will be how the entire semiconductor supply chain adapts and collaborates to support this AI-driven growth trajectory, ensuring sustained technological advancement.

Source: https://finance.biggo.com/news/d81ffb54-2ad9-4684-805d-14a08d4f7eeb

Collected: September 18, 2026 | Automated Research System (Gemini API)

#07 Google Reportedly Considering CoWoS as Backup for 2027 TPUs Amid Intel EMIB-T Substrate Yield Lag Reports

Published September 17, 2026　(情報元の名称なし)　USA

OVERVIEW

Google's Humufish TPU project, slated for mass production by late 2027, is reportedly considering TSMC's CoWoS as a backup packaging solution, following reports of lagging substrate yields for Intel's EMIB-T technology. While ASIC partner MediaTek expressed confidence in EMIB-T yields for stable mass production, Google's dual-sourcing strategy highlights the complexities of advanced packaging and the critical importance of risk management. This move underscores major tech firms' strategies to stabilize the supply of AI accelerators.

IN DEPTH

Key Findings

Google's next-generation Humufish TPU (Tensor Processing Unit) project, initially planned to utilize Intel's EMIB-T advanced packaging technology for mass production by late 2027, is now reportedly considering TSMC's CoWoS (Chip-on-Wafer-on-Substrate) as a backup solution. This consideration arises from reports of lagging substrate yields for EMIB-T, indicating major tech companies' strategic risk management in securing a stable supply of AI accelerators.

Technical / Clinical Details

The Google Humufish TPU project, supported by ASIC (Application-Specific Integrated Circuit) partner MediaTek, aims for mass production by the end of 2027 using Intel's EMIB-T (Embedded Multi-die Interconnect Bridge - Tiled) packaging technology. EMIB-T is a crucial method for efficiently connecting multiple dies (chips) on a single package substrate, standing alongside CoWoS as a vital advanced packaging solution for enhancing the performance of AI and High-Performance Computing (HPC) chips. MediaTek, in its earnings call, expressed confidence that EMIB-T yields had reached a level capable of supporting stable mass production, indicating their readiness for the 2027 production schedule. However, recent reports suggest challenges in the EMIB-T substrate manufacturing process, compelling Google to explore CoWoS as a proven alternative to ensure supply stability.

Background & Context

The increasing complexity of AI models and the explosion of data volumes have dramatically escalated the demand for high-performance AI accelerators. Consequently, advanced packaging technologies have become a significant bottleneck in the semiconductor supply chain. Hyperscalers such as Google, Amazon, and Microsoft are investing heavily in their custom AI chips, and the stability of their supply directly impacts their data center expansion strategies. Both EMIB-T and CoWoS are powerful technologies for integrating multiple chips, but each possesses unique manufacturing challenges and advantages. Google's consideration of CoWoS as a backup likely stems from a strategic intent to avoid over-reliance on a single technology and to diversify supply chain risks.

Strategic Significance & Outlook

Google's strategic maneuver highlights the intense competition and inherent challenges within advanced packaging technologies. For Intel, improving EMIB-T yields will be an urgent priority to establish credibility in the custom AI chip market and compete effectively against CoWoS. Conversely, this situation could further increase demand for TSMC's CoWoS, making its ongoing capacity expansions even more critical. This dual-sourcing strategy demonstrates that securing flexibility and resilience in AI accelerator supply will be paramount for leading technology companies in the coming years, influencing both technological development competition and supply chain strategies across the semiconductor industry.

Source: https://www.digitimes.com/news/a20260917PD208/google-tpu-cowos-substrate-2027.html

Collected: September 18, 2026 | Automated Research System (Gemini API)

#08 CSIS Report: AI and HBM Demand Drive New Semiconductor Shortage, Highlighting Structural Shift and Supply Chain Security Concerns

Published September 17, 2026 CSIS USA

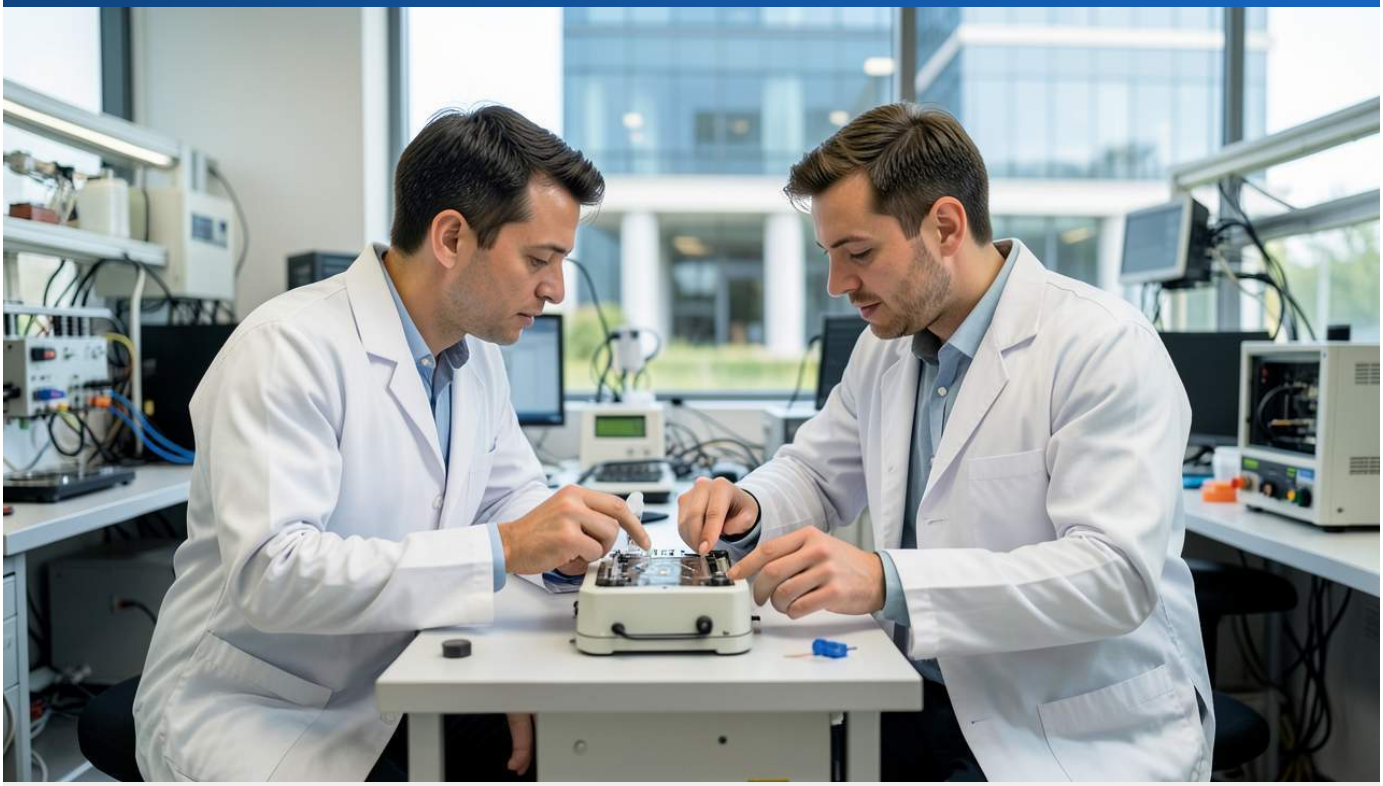

OVERVIEW

A CSIS report indicates that the current memory shortage reflects both structural demand growth from AI and cyclical market fluctuations. The rapid expansion of AI data centers has led to a surge in demand for HBM (High Bandwidth Memory), essential for AI workloads, which in turn strains conventional DRAM supply due to HBM's need for advanced packaging. The concentration of commercial-scale HBM production in Asia raises supply chain concentration risks and national security concerns, even with increasing U.S. investments, pointing to a new structural challenge in the semiconductor market.

IN DEPTH

Key Findings

A report from the Center for Strategic and International Studies (CSIS) reveals that the current semiconductor memory shortage is attributable to both a structural increase in demand driven by AI applications and typical cyclical market fluctuations. Critically, the rapid expansion of AI data centers has dramatically increased demand for High Bandwidth Memory (HBM), which is indispensable for AI workloads, and this has exacerbated the supply of traditional DRAM. The concentration of commercial-scale HBM production in Asia, particularly in South Korea and Taiwan, presents significant supply chain risks and national security concerns, even as investments in U.S. domestic semiconductor manufacturing advance.

Technical / Clinical Details

HBM is a specialized memory technology vital for AI workloads, offering superior bandwidth and lower power consumption crucial for maximizing the performance of AI accelerators and high-performance GPUs. While HBM is fabricated using standard DRAM, it necessitates advanced packaging capabilities (e.g., CoWoS) to stack multiple DRAM dies and integrate them with other chiplets. This additional, complex process not only constrains HBM supply but also strains the overall supply of conventional DRAM. The report emphasizes the geographical concentration of commercial HBM production in Asia, primarily South Korea and Taiwan. Despite ongoing investments in semiconductor manufacturing within the United States, this geographical centralization poses a potential single point of failure and raises significant national security concerns, underscoring the fragility of critical supply chains.

Background & Context

Historically, the memory market has been characterized by cyclical supply-demand fluctuations, often referred to as the "memory cycle." However, the advent of AI has introduced a new, structural demand for HBM. This fundamental shift is recognized as a long-term trend, not merely a temporary shortage, driven by the continuous evolution and proliferation of AI technologies. The fact that leading AI companies such as NVIDIA, Google, AMD, and Amazon consume the majority of the HBM supply illustrates its immense importance to the entire AI ecosystem. Simultaneously, in an era of heightened geopolitical risks, the concentration of critical semiconductor technology production in specific regions poses an urgent challenge for governments worldwide concerning supply chain resilience and national security.

Strategic Significance & Outlook

The HBM supply shortage and the geographical concentration of its supply chain are critical issues impacting both AI infrastructure expansion and national security strategies. Moving forward, nations, including the United States, are likely to pursue strategies aimed at localizing HBM manufacturing capabilities or diversifying production to friendly regions to enhance supply chain resilience. Additionally, accelerating technological innovation in HBM production to establish more efficient and cost-competitive manufacturing methods will be essential. Addressing these structural challenges is indispensable for the sustainable development of AI technologies and the stabilization of the global semiconductor supply chain, ensuring that the foundational components for future innovation are secure and robust.

Source: https://www.csis.org/analysis/beyond-memory-cycle-ai-hbm-and-new-semiconductor-shortage

Collected: September 18, 2026 | Automated Research System (Gemini API)

#09 ASE's SPIL Nears Arizona Back-End Test Facility Expansion Amidst $10.5B CAPEX Surge; Amkor Boosts US Packaging Investment to $12B

Published September 15, 2026 TrendForce Taiwan

OVERVIEW

ASE subsidiary SPIL is finalizing plans for a new back-end test facility in Arizona, part of ASE's $10.5 billion CAPEX for 2026, which includes $4 billion for new facilities and $6.5 billion for equipment. Concurrently, SPIL is expanding CoWoS advanced packaging in Douliu, Taiwan. Amkor Technology is also significantly increasing its Arizona investment from $7 billion to $12 billion, targeting production by 2028 and full capacity by 2030, highlighting a major push by OSAT leaders to bolster advanced packaging capabilities in the US.

IN DEPTH

Key Findings

SPIL, a subsidiary of Taiwanese semiconductor assembly and test giant ASE, is reportedly in the final stages of planning a new back-end test facility in Arizona, reinforcing the global semiconductor supply chain in North America. This expansion aligns with ASE's projected capital expenditure of $10.5 billion for 2026, with $4 billion allocated for new plants and facilities and $6.5 billion for equipment. This strategic move responds to increasing geopolitical risks and incentives provided by the U.S. government to onshore semiconductor manufacturing capabilities.

Technical / Clinical Details

The new SPIL facility in Arizona will primarily provide back-end test services, strategically locating operations closer to key U.S. customers. In parallel, SPIL is aggressively expanding its advanced packaging capabilities in Taiwan, with construction for a new factory in Douliu commencing in mid-August. This facility will integrate advanced CoWoS (Chip-on-Wafer-on-Substrate) packaging technology, crucial for meeting the surging demand for high-performance semiconductors, particularly AI chips.

Furthermore, industry peer Amkor Technology has substantially increased its investment in Arizona from an initial $7 billion to $12 billion. Amkor's expanded operations are slated to begin production by 2028 and reach full capacity by 2030. These significant investments underscore the commitment of leading OSAT (Outsourced Semiconductor Assembly and Test) providers to build out a robust advanced semiconductor packaging ecosystem within the United States.

Background & Context

As the limits of semiconductor process node scaling are approached, advanced packaging technologies—including chiplet integration, 3D stacking, and heterogeneous integration—are becoming pivotal for enhancing device functionality. Demand in high-growth sectors such as AI, High-Performance Computing (HPC), and automotive electronics is rapidly accelerating the adoption of these technologies. The U.S., through initiatives like the CHIPS Act, is actively promoting domestic manufacturing, offering incentives to OSAT companies to establish and expand facilities within its borders.

Strategic Significance & Outlook

The substantial back-end investments by ASE and Amkor in the U.S., in conjunction with the construction of major foundry facilities by companies like TSMC, represent a critical step towards realizing a comprehensive U.S.-based semiconductor supply chain. This not only mitigates geopolitical supply chain risks but also promises to accelerate technological innovation and bolster domestic production capacity for high-value semiconductor products. While intensifying competition among OSAT providers, sustained investment in advanced packaging technologies will continue to drive overall industry growth and evolution.

Source: https://www.trendforce.com/news/2026/09/15/news-ases-spil-reportedly-nears-arizona-expansion-as-tsmc-amkor-step-up-u-s-advanced-packaging-push/

Collected: September 18, 2026 | Automated Research System (Gemini API)

#10 Samsung's zHBM Concept Promises 70% Lower I/O Power, 2.3x Bandwidth vs. HBM5 via Hybrid Bonding for Edge AI

Published September 16, 2026 Tom's Hardware South Korea

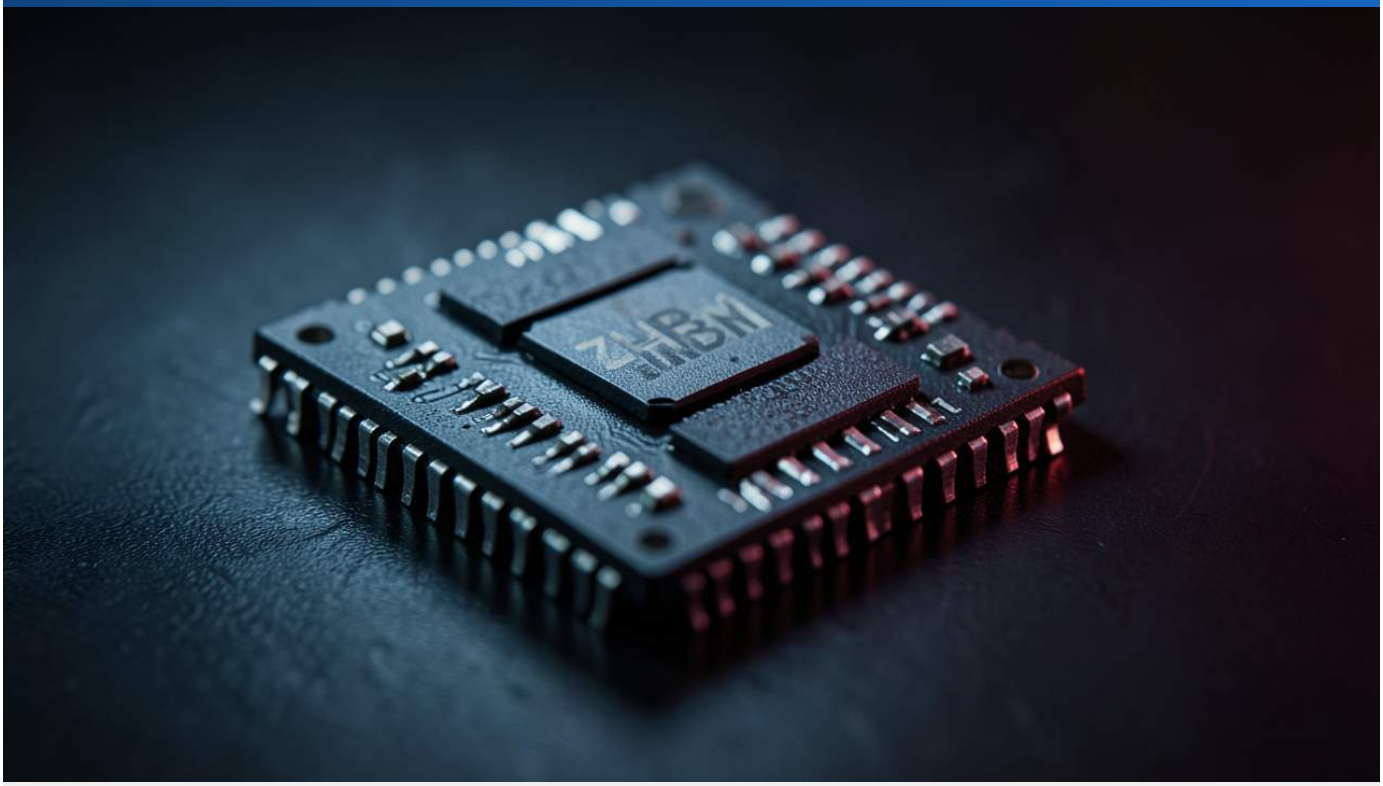

OVERVIEW

Samsung proposes zHBM, an innovative memory architecture that directly stacks DRAM onto the processor, bypassing the interposer, and is projected to reduce I/O power by approximately 70% compared to HBM5 while offering 2.3 times the bandwidth of HBM4E. This technology necessitates wafer-on-wafer hybrid copper bonding and tight co-design between DRAM and SoC teams. An SK hynix executive suggests hybrid bonding may not be ready for HBM4E, anticipating HBM5 as the earliest adoption point, signaling a major shift for high-performance, low-power memory solutions in edge AI.

IN DEPTH

Key Findings

Samsung is pioneering a concept called 'zHBM,' a next-generation memory technology poised to revolutionize high-bandwidth memory (HBM) architectures by directly stacking DRAM onto the processor rather than relying on an interposer. This innovative approach is projected to achieve approximately 70% lower I/O power consumption compared to HBM5 and deliver 2.3 times the bandwidth of HBM4E. Such an advancement promises significant improvements in performance and power efficiency, particularly for edge AI devices.

Technical / Clinical Details

The core of the zHBM technology lies in wafer-on-wafer hybrid copper bonding. This method enables direct, fine-pitch electrical connections between chips, offering higher density and superior electrical characteristics than traditional thermocompression bonding with microbumps. However, realizing zHBM requires unprecedentedly close co-design efforts between DRAM and System-on-Chip (SoC) teams, elevating the level of memory-logic integration. Jaesik Lee, VP of Packaging Engineering at SK hynix, commented that hybrid bonding technology is not yet mature for the HBM4E generation, anticipating its earliest viable adoption will be with HBM5, which necessitates 20-layer DRAM stacking.

Background & Context

The exponential growth in AI processing demands has made memory bandwidth and power efficiency critical bottlenecks, especially in edge AI devices and data centers. While HBM emerged to address these challenges, the continuous push for higher performance and lower power consumption necessitates further innovation. Conventional HBM connects to processors via an interposer, leading to longer signal paths, power loss, and latency. zHBM seeks to fundamentally resolve these issues by eliminating the interposer.

Strategic Significance & Outlook

Direct stacking technologies, with hybrid bonding at their core, like zHBM, have the potential to profoundly impact future edge AI chip designs, enabling the development of smaller, more efficient, and higher-performing AI devices. However, achieving high-precision wafer-level bonding, fostering close collaboration across different vendors, and establishing new design and manufacturing processes will require several more years of development. Anticipated for adoption with the HBM5 generation, this technology represents a crucial milestone that will further drive innovation in the semiconductor back-end industry.

Source: https://www.tomshardware.com/tech-industry/semiconductors/piecemakers-bets-edge-ai-devices-will-diverge-from-reliance-on-hbm-custom-designed-memory-fuses-dram-stack-directly-to-the-processor-using-hybrid-bonding

Collected: September 18, 2026 | Automated Research System (Gemini API)

#11 BESI Hybrid Bonding System Orders Exceed 150 Units, HBM5 Onward to Accelerate Memory Adoption

Published September 17, 2026 PhotonCap Unknown

OVERVIEW

BESI has secured over 150 cumulative orders for its hybrid bonding systems, with six integrated production lines (including 30 BESI bonders) implemented by major logic customers in partnership with Applied Materials. Currently deployed primarily for logic and Co-packaged Optics (CPO), this technology is expected to see memory adoption, specifically for HBM5 (20-layer stacking) and beyond, after HBM4E continues with existing methods, anticipating a three-year delay in memory sector penetration. This trend underscores the increasing criticality of hybrid bonding in advanced packaging.

IN DEPTH

Key Findings

BE Semiconductor Industries (BESI) has reported over 150 cumulative orders for its innovative hybrid bonding systems, signaling robust market demand and expanding adoption within advanced packaging technologies. Notably, six integrated production lines, developed in collaboration with Applied Materials and incorporating 30 BESI bonders, have been deployed at key logic semiconductor customers.

Technical / Clinical Details

Hybrid bonding is a direct copper-to-copper bonding technology that operates without heat or pressure, enabling ultra-fine pitch connections and superior signal transmission characteristics. This technique is currently vital for manufacturing high-performance logic chips and Co-packaged Optics (CPO) devices, playing a crucial role in chiplet technology and 3D stacking. CPO, by integrating optical modules and semiconductors into a single package, promises significant improvements in data center bandwidth and power efficiency.

However, the adoption of hybrid bonding in the memory sector is occurring with a time lag. For High Bandwidth Memory (HBM), existing methods like wire bonding and micro-bump-based flip-chip bonding are projected to remain dominant until the HBM4E generation. Full-scale adoption of hybrid bonding for memory is anticipated from HBM5 (20-layer stacking) onward, with an estimated three-year delay compared to logic applications. This delay is attributed to the complex stacking structures and manufacturing processes inherent in memory chips.

Background & Context

The explosive demand for AI and High-Performance Computing (HPC) necessitates continuous advancements in semiconductor chip performance and miniaturization. To meet these demands, advanced packaging technologies, which integrate chips vertically and horizontally, have become a primary driving force in the semiconductor industry. Hybrid bonding is gaining prominence as a key enabler for next-generation 3D-ICs and chiplet integration, offering higher connection density and superior electrical performance compared to conventional thermocompression bonding.

Strategic Significance & Outlook

BESI's strong order book for hybrid bonding systems confirms the technology's indispensable role in future semiconductor manufacturing. Despite the delay in memory applications, with HBM5 onwards projected for adoption, hybrid bonding is expected to expand its reach across logic, CPO, and memory segments in the coming years. This will enable semiconductor manufacturers to achieve further performance enhancements and form factor reductions, thereby accelerating the evolution of the AI and HPC markets.

Source: https://photoncap.net/p/investment-map-22-listed-companies

Collected: September 18, 2026 | Automated Research System (Gemini API)

#12 AI Accelerator ABF Substrate Shortage to Exceed 40% by 2028; Ibiden, Unimicron, Samsung Electro-Mechanics Rush Billions in Capacity Expansion

Published September 10, 2026 Tom's Hardware Japan

OVERVIEW

The supply shortage of ABF (Ajinomoto Build-up Film) substrates for AI accelerators is projected to worsen significantly, potentially reaching approximately 10% by late 2026, 21% by 2027, and over 40% by 2028. In response, Ibiden is investing ¥500 billion (approx. $3.1B) from 2026-2028 to increase production capacity 2.8-fold compared to 2024. Unimicron is dedicating a record NT$34 billion (approx. $1.07B) to CAPEX in 2026, while Samsung Electro-Mechanics is investing $1.2 billion in ABF substrate production, aiming for mass production by Q3 2027. Ajinomoto, controlling over 95% of the ABF film market, was already at full capacity in Q2 2026 and announced a roughly 30% price increase in May.

IN DEPTH

Key Findings

As the demand for AI accelerators skyrockets, a critical shortage of indispensable ABF (Ajinomoto Build-up Film) substrates is casting a significant shadow over the industry. Current projections indicate that the ABF substrate supply deficit could reach approximately 10% by late 2026, escalate to about 21% by 2027, and potentially surpass 40% by 2028. In response, major suppliers such as Ibiden, Unimicron, and Samsung Electro-Mechanics are aggressively planning and executing multi-billion dollar capital investments to rapidly expand their production capacities.

Technical / Clinical Details

ABF substrates are crucial packaging materials that enable high-speed signal transmission and high-density wiring in high-performance AI chips and server processors. Their scarcity directly impedes the scaling of AI semiconductor production. In this critical situation, companies are advancing concrete expansion plans:

- **Ibiden (Japan):** Plans to invest ¥500 billion (approximately $3.1 billion) between 2026 and 2028, aiming to expand its production capacity for AI server substrates by 2.8 times compared to 2024 levels.
- **Unimicron (Taiwan):** Is increasing its capital expenditure for 2026 to a record NT$34 billion (approximately $1.07 billion), primarily focusing on ABF substrate production.
- **Samsung Electro-Mechanics (South Korea):** Is investing $1.2 billion in ABF substrate production, with mass production targeted to commence by the third quarter of 2027.

Furthermore, Ajinomoto (Japan), which commands over 95% of the ABF film market, was already operating at full capacity in Q2 2026. Due to severe demand pressure, the company notified customers of an approximate 30% price increase in May of the same year, indicating that raw material supply constraints and rising costs are propagating throughout the entire supply chain.

Background & Context

The increasing complexity and scale of AI models are continually raising performance requirements for AI accelerators. This necessitates integration with High Bandwidth Memory (HBM) and larger chip designs to accommodate more transistors, consequently demanding larger and more complex ABF substrates. However, ABF substrate manufacturing requires advanced technology, specialized equipment, and long lead times, making it challenging to respond to the rapid surge in demand. A key factor contributing to this bottleneck is the concentrated supply of ABF film, with Ajinomoto being virtually the sole supplier.

Strategic Significance & Outlook

While the substantial capital investments by leading manufacturers are expected to alleviate the long-term supply shortage, it will take time for these investments to translate into actual market capacity. The severe supply deficit projected through 2028 could significantly impact the growth trajectory of the entire AI semiconductor industry. A stable supply of ABF substrates is crucial for building next-generation AI infrastructure, necessitating further collaboration and investment to enhance overall supply chain resilience. Additionally, diversification of material suppliers and the development of alternative materials will be critical challenges moving forward.

Source: https://www.tomshardware.com/tech-industry/semiconductors/the-state-of-abf-substrates-in-data-center-silicon-in-2026-solving-the-supply-crunch-and-material-wall-beneath-every-ai-accelerator

Collected: September 18, 2026 | Automated Research System (Gemini API)

#13 Samsung Electro-Mechanics Unveils Next-Gen Package Substrate Technologies for AI & Servers at KPCA Show 2026, Highlighting Warp Reduction & High-Density Interconnects

Published September 10, 2026 Samsung Electro-Mechanics South Korea

OVERVIEW

Samsung Electro-Mechanics showcased a comprehensive suite of next-generation semiconductor package substrate technologies targeting the AI and server markets at 'KPCA Show 2026.' The company presented diverse advanced packaging solutions, including 2.5D and 2.1D package substrates, glass substrates, automotive FCBGAs, and Ultra Thin Chip Scale Package (UTCSP) substrates. Emphasis was placed on core technologies for AI accelerator 2.5D package substrates, such as warp reduction for large, multi-layer substrates, high-speed signal transmission, and high-density interconnection.

IN DEPTH

Key Findings

Samsung Electro-Mechanics unveiled a comprehensive portfolio of next-generation semiconductor package substrate technologies optimized for AI and server applications at the 'KPCA Show 2026.' The company showcased innovative solutions tailored for a broad range of high-performance sectors, including AI, servers, data centers, mobile, and autonomous driving. Key advancements highlighted include warp reduction technology for large, multi-layer substrates, ultra-high-speed signal transmission capabilities, and high-density interconnection techniques, all critical for their 2.5D package substrates targeting AI accelerators.

Technical / Clinical Details

The technologies presented by Samsung Electro-Mechanics are designed to meet the demanding requirements of today's high-performance semiconductors. Their 2.5D package substrates for AI accelerators integrate multiple chips (e.g., GPUs and HBMs) on an interposer to achieve high bandwidth and low latency. A critical aspect here is the effective management of warp issues, which become more pronounced with larger and multi-layer substrates. The company addresses this challenge through proprietary materials and structural designs, ensuring reliable packaging.

Furthermore, 2.1D package substrates are specialized for applications demanding a balance between cost-efficiency and performance. Glass substrates, an emerging technology for future packaging, offer superior flatness, thermal stability, and electrical properties compared to conventional organic substrates, enabling finer wiring and higher integration. Automotive FCBGAs are designed to meet the high reliability and thermal performance required by autonomous driving systems, while UTCSP substrates are positioned as ultra-thin, high-density packaging solutions for mobile devices.

Background & Context

Innovations in fields such as AI, data centers, 5G communication, and autonomous driving are placing unprecedented demands on semiconductor packaging technologies. The pursuit of higher performance, greater functionality, and miniaturization in chips necessitates package substrates with increased density, faster signal transmission rates, and enhanced thermal management capabilities. Specifically for AI accelerators, an extremely high-bandwidth connection between GPUs and HBMs is essential for processing vast amounts of data, making 2.5D packaging a mainstream solution. The introduction of new materials like glass substrates holds the potential to further push these performance boundaries.

Strategic Significance & Outlook

Samsung Electro-Mechanics' exhibition underscores its position as a leading innovator in the advanced semiconductor package substrate market. By offering a wide array of solutions for high-growth sectors like AI, servers, and automotive, the company is well-positioned to capture significant demand from these markets. Critically, investments in next-generation technologies such as glass substrates represent a strategic move to address future needs for ultra-high-density and high-performance packaging, contributing to the overall technological evolution of the semiconductor industry.

Source: https://www.semiconductor-digest.com/samsung-electro-mechanics-showcases-next-generation-semiconductor-package-substrate-technologies-for-ai-and-servers-at-kpca-show-2026/

#14 AI Infrastructure Component Crunch Becomes Reality: ABF Substrate Lead Times Hit 56 Weeks, Equipment Parts Up to 40 Months

Published September 16, 2026 BigGo Finance Taiwan

OVERVIEW

The supply shortage for AI infrastructure components is becoming entrenched; TrendForce reports ABF substrate lead times surged to 48–56 weeks, approximately 4 to 4.7 times the equilibrium. HDDs reached 50 weeks, DRAM 20 weeks, and MLCCs 30 weeks, with all five tracked categories (excluding GPUs) in shortage. Semiconductor equipment parts are particularly critical, with some Japanese-made crucial components facing lead times up to 40 months. Structural factors like slow component supplier expansion, rising CAPEX, and HBM production burdens make short-term resolution unlikely.

IN DEPTH

Key Findings

The supply shortage for critical semiconductor components essential for AI infrastructure is intensifying, transitioning from a temporary issue to a structural crisis. According to the latest TrendForce analysis, lead times for ABF (Ajinomoto Build-up Film) substrates, a key material for AI accelerators, have skyrocketed to 48–56 weeks, representing approximately 4 to 4.7 times the normal equilibrium. Furthermore, other major components such as Hard Disk Drives (HDDs) at 50 weeks, DRAM at 20 weeks, and Multi-Layer Ceramic Capacitors (MLCCs) at 30 weeks are also experiencing significant shortages, with all five major categories tracked (excluding GPUs) now in deficit. Alarmingly, lead times for some critical semiconductor manufacturing equipment parts, particularly those from Japan, are reportedly extending up to an unprecedented 40 months.

Technical / Clinical Details

ABF substrates serve as vital interconnect foundations for integrating multiple chips (CPUs, GPUs, HBMs, etc.) in high-performance AI processors. Their manufacturing demands sophisticated process technology and dedicated equipment, inherently leading to long lead times. The current dramatic increase in lead times underscores the demand for AI significantly outpacing forecasts and the limits of existing production capacities.

HDDs, DRAM, and MLCCs are also indispensable for building AI data centers and edge devices. The scarcity of these components will delay the overall deployment of AI systems. The most critical aspect is the supply issue for semiconductor manufacturing equipment parts. As lead times for components required by equipment manufacturers (e.g., precision-machined parts, optical components, control systems) extend, it creates a vicious cycle where the expansion of semiconductor manufacturing capacity itself is hindered.

This supply crunch stems from multiple structural factors: the slow pace of capacity expansion by component manufacturers, rising raw material prices and capital expenditure costs, and the technical and resource burden associated with High Bandwidth Memory (HBM) production. These intertwined factors make a short-term resolution extremely challenging.

Background & Context

The rapid advancement of AI technology is accelerating the expansion of data center scale and processing power, consequently driving an exponential increase in demand for related semiconductor and electronic components. However, the semiconductor supply chain is an extremely complex and multi-layered structure, possessing vulnerabilities where bottlenecks in one area can cascade throughout the entire system. Specifically, ABF substrates, which require advanced manufacturing processes, and manufacturing equipment parts that depend on specific suppliers, inherently struggle to keep pace with rapid demand surges. Unlike previous semiconductor shortages, the current situation is characterized by structural and composite factors rather than a single market driver.

Strategic Significance & Outlook

The normalization of AI infrastructure component shortages is likely to have long-term implications for the growth trajectory of the AI industry. Semiconductor manufacturers and system integrators will be compelled to prioritize supply chain resilience, accelerating multi-supplier strategies, optimizing inventory management, and investing in alternative technologies. Policymakers are also urged to further strengthen international cooperation and support for domestic manufacturing capabilities to alleviate bottlenecks across the entire semiconductor ecosystem. The prolonged lead times for manufacturing equipment parts pose a risk to the overall semiconductor industry's capital expenditure cycle, potentially further delaying the expansion of future supply capacity.

Source: https://finance.biggo.com/news/02cf5c8f-c58d-45f8-a7a3-8fe37716bc8f

Collected: September 18, 2026 | Automated Research System (Gemini API)

#15 Glass Substrates & TGV Technology Set to Reshape 2027 Advanced Packaging Race; US, Japan, Korea, Taiwan, China Accelerate Development

Published September 16, 2026 digitimes Taiwan

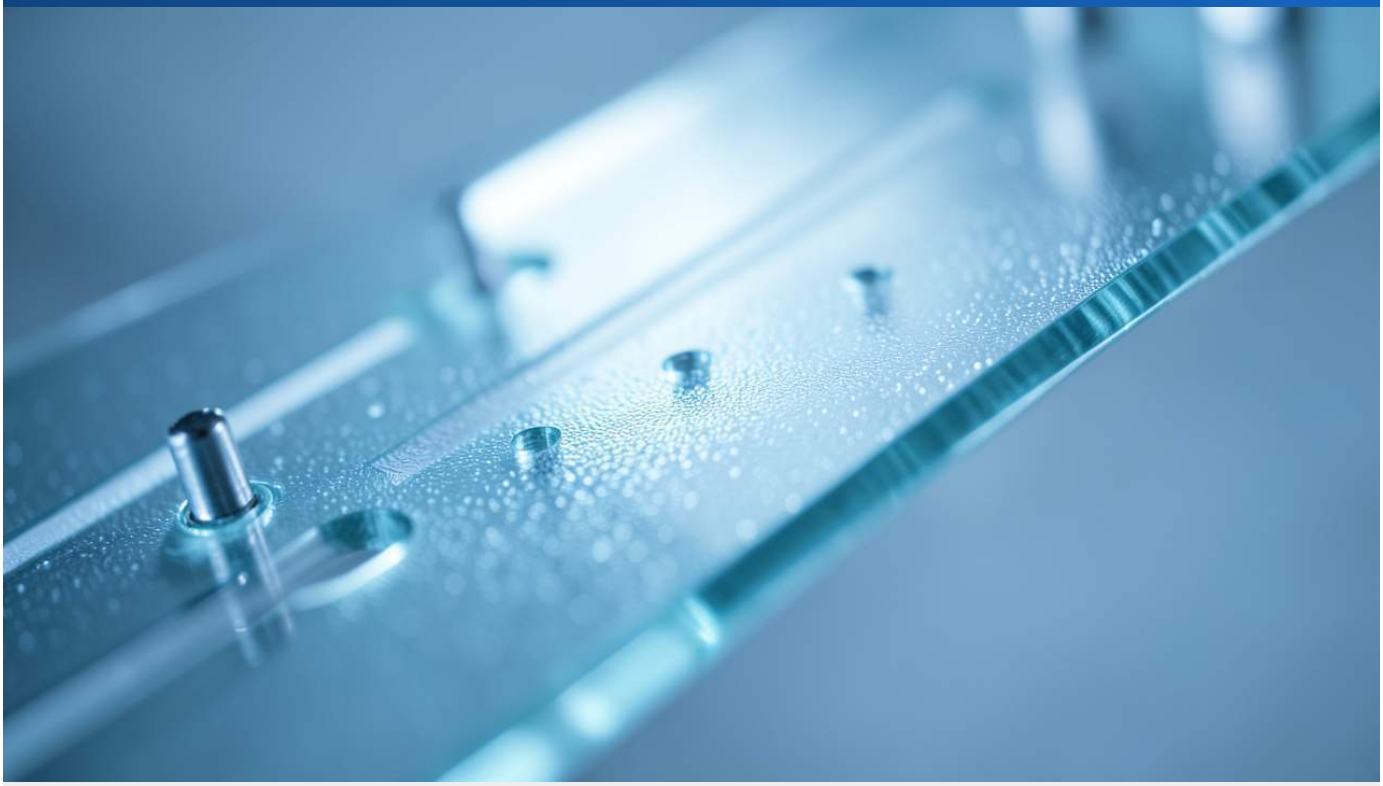

OVERVIEW

As AI chip package sizes expand, glass substrates and Through-Glass Via (TGV) technology are emerging as new focal points in advanced packaging competition. Supply chain leaders in the US, Japan, South Korea, Taiwan, and China are accelerating development and preparation for market opportunities in 2027 and 2028, focusing on four key steps: glass material development, via formation, via metallization, and multi-layer substrate lamination. This technology is expected to deliver performance improvements and cost reductions beyond the limits of conventional organic substrates.

IN DEPTH

Key Findings

With the increasing package sizes of AI chips, glass substrates and Through-Glass Via (TGV) technology are rapidly emerging as the next frontier in advanced packaging. This technology holds the potential for performance enhancements and cost reductions beyond the limitations of conventional organic substrates. Major supply chain companies across the United States, Japan, South Korea, Taiwan, and China are accelerating their development and mass production readiness, targeting market opportunities in 2027 and 2028.

Technical / Clinical Details

Glass substrates offer superior flatness, lower coefficient of thermal expansion (CTE), mechanical strength, and low dielectric loss compared to organic substrates. These characteristics enable finer wiring patterns and ultra-high-density interconnects. TGV technology specifically involves forming vertical vias through the glass substrate and filling them with metal to establish electrical connections. This technique shortens signal transmission distances, improves power efficiency, and allows for higher-layer 3D stacking, making it essential for meeting the stringent requirements of AI accelerators and High-Performance Computing (HPC) chips.

Development of glass substrate and TGV technology is focused on four key steps:

- **Glass Material Development:** Selecting and manufacturing glass materials with suitable properties for packaging applications.
- **Via Formation Technology:** High-precision TGV creation through methods like laser processing or etching.
- **Via Metallization:** The process of filling the formed vias with metals such as copper.
- **Multi-layer Substrate Lamination:** Processes for stacking glass substrates integrated with TGVs to form multi-layered structures.

In each of these steps, materials suppliers in the U.S., equipment manufacturers in Japan, OSATs in Korea, foundries in Taiwan, and packaging companies in China are leveraging their respective strengths in an intense development race.

Background & Context

Current organic package substrates are beginning to face physical limitations in meeting the demands for thermal management, signal integrity, and larger form factors driven by AI chip performance advancements. Particularly in 2.5D/3D packaging, which integrates HBM and logic chips, the need for larger areas, more layers, and finer pitches exacerbates issues like warpage and transmission loss in organic substrates. Glass substrates are anticipated to resolve these challenges, positioning them as a critical next-generation platform in the overall semiconductor industry roadmap.

Strategic Significance & Outlook

The introduction of glass substrates and TGV technology has the potential to significantly transform the competitive landscape of advanced packaging from 2027 onwards. Once this technology matures and mass production capabilities are established, it is expected to further enhance the performance and power efficiency of AI chips while also contributing to cost optimization. Aggressive investments and collaborations by leading nations and companies are anticipated to accelerate the practical application of glass substrate technology, strongly supporting the development of next-generation AI and HPC systems. This technology is poised to be a key driver for a new innovation cycle within the semiconductor industry.

Source: https://www.digitimes.com/news/a20260916PD221/ai-chip-2027-substrate-technology-materials.html

Collected: September 18, 2026 | Automated Research System (Gemini API)

#16 Samsung Electro-Mechanics and Qualcomm Jointly Develop Organic Bridge for Advanced 2.1D AI Semiconductor Packaging

Published September 14, 2026 The Elec South Korea

OVERVIEW

Samsung Electro-Mechanics and Qualcomm have collaborated to develop an 'organic bridge' technology for advanced 2.1D packaging in AI semiconductors. This innovation facilitates high-density, high-bandwidth interconnects between dies like GPUs and HBMs by utilizing organic materials in the substrate manufacturing process, differing from Intel's silicon-based EMIB. This approach aims to boost AI chip performance while offering a potentially more cost-effective and flexible packaging solution.

IN DEPTH

Key Findings

Samsung Electro-Mechanics, in collaboration with Qualcomm, has successfully co-developed an 'organic bridge' technology for advanced 2.1D packaging, specifically targeting AI semiconductors. This breakthrough enables significantly increased input/output (I/O) density and bandwidth between disparate dies, such as graphics processing units (GPUs) and high-bandwidth memory (HBM), which is critical for enhancing AI chip computational capabilities. The organic bridge serves a similar function to Intel's silicon-based EMIB but leverages organic materials within the substrate manufacturing process rather than traditional semiconductor fabrication, offering a novel approach to multi-die integration.

Technical Details

The organic bridge technology distinguishes itself by employing organic materials manufactured during the substrate production process, presenting a more flexible and potentially cost-efficient alternative to silicon-based interposers. This method effectively creates high-density interconnects, allowing for the seamless integration of multiple chips within a single package. By utilizing organic substrates, the technology aims to circumvent some of the manufacturing complexities and material costs associated with silicon interposers, while still achieving the necessary electrical performance for advanced packaging. This allows for closer proximity and faster communication between logic and memory components, a bottleneck in many high-performance computing applications.

Background & Context

The accelerating demand for AI computing power has pushed the boundaries of traditional chip design, making advanced packaging a critical factor in semiconductor performance. The integration of various functional dies, like HBM and high-performance processors, into a single package (2.5D/3D packaging) is essential, but current silicon-based interposer solutions face challenges in cost, yield, and manufacturing complexity. The organic bridge emerges as a strategic alternative, particularly in the 2.1D packaging domain, which bridges the gap between traditional 2D and more complex 2.5D structures. This development by Samsung Electro-Mechanics and Qualcomm signals a growing trend towards material diversification and process innovation in the advanced packaging landscape, offering chip designers more options to optimize performance and cost for AI accelerators.

Strategic Significance & Outlook

The introduction of organic bridge technology holds significant implications for the future of AI semiconductor development. It promises to enable higher performance at a potentially lower cost, fostering broader adoption of advanced packaging across various AI applications, from data centers to edge devices. This joint development highlights the industry's continuous effort to innovate beyond silicon for inter-die connectivity. The collaboration between a leading component manufacturer and a major chip designer could accelerate the standardization and commercialization of such organic-based packaging solutions, potentially reshaping the competitive landscape of advanced packaging and contributing to the overall advancement of the AI ecosystem.

Source: https://www.thelec.net/news/articleView.html?idxno=13907

Collected: September 18, 2026 | Automated Research System (Gemini API)