

TROY TECHNICAL

SEMICONDUCTOR BACK-END WEEKLY REPORT

October 3 - October 9, 2026 | Issue 2026-W40

**Package performance, process yield, test
coverage and qualified capacity**


15

ANALYZED ARTICLES

10

PRIORITY THEMES

Issued 2026-10-09 | English edition | fixed approved cover artwork



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

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

SIGNAL 01 | MATERIALS

GlobalFoundries: \$2B TSMC CoWoS interposer US production timeline

GlobalFoundries and TSMC have signed a 5-year, \$2 billion agreement for GlobalFoundries to manufacture silicon interposers for TSMC's CoWoS advanced packaging ecosystem.

SIGNAL 02 | BONDING & ASSEMBLY

Towa: \$31.6M AI chip packaging factory and 2029 supply specs

Japanese semiconductor equipment manufacturer Towa announced plans to invest approximately ¥5 billion (\$31.6 million) in a new factory in Kyoto Prefecture to meet increasing demand for advanced semiconductor packaging equipment driven by AI infrastructure.

SIGNAL 03 | ADVANCED PACKAGE

AMD: Taiwan packaging and substrate capacity for 2027/2028

AMD CEO Lisa Su visited Taiwan to secure advanced process technology, packaging, and substrate capacity for 2027 and 2028, ensuring smooth shipments of next-generation AI server products.

SIGNAL 04 | TEST & INSPECTION

Glass Substrates: Intel Clearwater Forest 2026 roadmap

In 2026, facing the physical limits of organic substrates, the semiconductor industry is widely adopting glass substrates for advanced packaging.

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

\$2 billion

A07

SOURCE-BOUND CONDITION

GlobalFoundries and TSMC have signed a 5-year, \$2 billion agreement for GlobalFoundries to manufacture silicon interposers for TSMC's CoWoS advanced packaging ecosystem.

5 billion

A13

SOURCE-BOUND CONDITION

Japanese semiconductor equipment manufacturer Towa announced plans to invest approximately ¥5 billion (\$31.6 million) in a new factory in Kyoto Prefecture to meet increasing demand for advanced semiconductor packaging.

2nm

A09

SOURCE-BOUND CONDITION

While TSMC accelerates 2nm and CoWoS-L expansion, the actual bottlenecks remain HBM memory supply, bit output degradation, and silicon interposer/package test yields.

36%

A14

SOURCE-BOUND CONDITION

Server CPU shipments from Intel and AMD are projected to increase by 36% and 62% respectively in 2027, driven by new platforms requiring 56%-63% larger substrate areas, further widening the ABF substrate gap.

4nm

A01

SOURCE-BOUND CONDITION

Samsung Electronics has developed an advanced HBM4E thermal management solution featuring a 4nm logic-based die and hybrid copper bonding (HCB), slashing thermal resistance by over 20% compared to HBM3E.

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.

1

Materials

P01/P06

GlobalFoundries and TSMC have signed a 5-year, \$2 billion agreement for GlobalFoundries to manufacture silicon interposers.

2

Bonding & assembly

P02/P07

Japanese semiconductor or equipment manufacturer Towa announced plans to invest approximately ¥5 billion (\$31.6 million) in a.

3

Advanced package

P03/P08

AMD CEO Lisa Su visited Taiwan to secure advanced process technology, packaging, and substrate capacity for 2027 and 2028.

4

Test & inspection

P04/P09

In 2026, facing the physical limits of organic substrates, the semiconductor industry is widely adopting glass substrates for.

5

Volume supply

P05/P10

Japanese PCB and substrate technology developer Elephantech announced it will open its new Toyosu Advanced Packaging Center.


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 next decision locus



REPRODUCIBLE POSITIONING RULE
Horizontal position uses five discrete stages: Research, Prototype, Joint evaluation, Product adoption and Scale. The vertical bands classify the next decision locus: Tech / process, System / adoption, or Business / rules (including supply systems). Article-title cues take precedence; the overview is used only when the title is silent. These are categories, not impact ranks. Bubble diameter reflects the editorial score.

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

P01

GlobalFoundries: \$2B TSMC CoWoS interposer US production timeline

A07 | Materials | Scale-up

Date not stated

FACT

GlobalFoundries and TSMC have signed a 5-year, \$2 billion agreement for GlobalFoundries to manufacture silicon interposers for TSMC's CoWoS advanced packaging ecosystem. This strategic partnership will significantly expand GlobalFoundries' production capacity at its Malta, New York facility, with mass production slated to begin in early 2028. The deal addresses a critical bottleneck in the AI semiconductor supply chain by establishing a domestic US source for these essential components.

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 5/5 | DECISION LOCUS: BUSINESS / RULES

P02

Towa: \$31.6M AI chip packaging factory and 2029 supply specs

A13 | Bonding & assembly | Product adoption

Date not stated

FACT

Japanese semiconductor equipment manufacturer Towa announced plans to invest approximately ¥5 billion (\$31.6 million) in a new factory in Kyoto Prefecture to meet increasing demand for advanced semiconductor packaging equipment driven by AI infrastructure expansion.

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 4/5 | DECISION LOCUS: BUSINESS / RULES

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

P03

AMD: Taiwan packaging and substrate capacity for 2027/2028

A09 | Advanced package | Product adoption

Date not stated

FACT

AMD CEO Lisa Su visited Taiwan to secure advanced process technology, packaging, and substrate capacity for 2027 and 2028, ensuring smooth shipments of next-generation AI server products. While TSMC accelerates 2nm and CoWoS-L expansion, the actual bottlenecks remain HBM memory supply, bit output degradation, and silicon interposer/package test yields. AMD's MI450 series AI platform, utilizing TSMC's 2nm and CoWoS, is slated for higher volume shipments in 2027.

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 4/5 | DECISION LOCUS: BUSINESS / RULES

P04

Glass Substrates: Intel Clearwater Forest 2026 roadmap

A12 | Test & inspection | Product adoption

Date not stated

FACT

In 2026, facing the physical limits of organic substrates, the semiconductor industry is widely adopting glass substrates for advanced packaging. Intel plans to ship its "Clearwater Forest" processor, the first with a glass core substrate, while TSMC will launch its CoWoS glass pilot line mid-year, and Samsung and Absolics are commencing new factory constructions.

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 20/25 | MATURITY 4/5 | DECISION LOCUS: SYSTEM / ADOPTION

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

P05

Elephantech: Toyosu Packaging Center 2027 launch and specs

A15 | Volume supply | Research

Date not stated

FACT

Japanese PCB and substrate technology developer Elephantech announced it will open its new Toyosu Advanced Packaging Center in Tokyo by March 2027. This center aims to significantly scale up manufacturing and R&D for high-performance PCBs and package substrates specifically designed for AI servers. Leveraging proprietary nanomaterials and precision inkjet processes, Elephantech seeks to meet the escalating demand for high-density interconnects required by next-generation AI infrastructure.

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 20/25 | MATURITY 1/5 | DECISION LOCUS: SYSTEM / ADOPTION

P06

ABF Substrates: 2026 production timeline and layer bottlenecks

A11 | Materials | Product adoption

Date not stated

FACT

As AI accelerators grow in size, ABF (Ajinomoto Build-up Film) substrates are emerging as a critical supply bottleneck in semiconductor back-end processing. Next-generation AI chips are expected to utilize ABF substrates with 18-20 layers by 2026, significantly increasing manufacturing steps like lamination, laser drilling, electroplating, etching, and alignment, thereby extending overall production time.

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 20/25 | MATURITY 4/5 | DECISION LOCUS: TECH / PROCESS

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

P07

Ibiden: 2027 server CPU substrate area and growth projections

A14 | Bonding & assembly | Product adoption

Date not stated

FACT

Ibiden is leveraging its market position due to a severe shortage of qualified high-end ABF substrates, which cannot keep pace with the increasing area and layer counts of AI silicon. Server CPU shipments from Intel and AMD are projected to increase by 36% and 62% respectively in 2027, driven by new platforms requiring 56%-63% larger substrate areas, further widening the ABF substrate gap.

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 | DECISION LOCUS: BUSINESS / RULES

P08

HBM4: SK hynix and Samsung 16-layer vs NVIDIA 12-layer specs

A08 | Advanced package | Product adoption

Date not stated

FACT

Despite NVIDIA's Vera Rubin AI platform favoring 12-layer HBM4 chips, SK hynix and Samsung Electronics are aggressively pursuing 16-layer HBM4 technology. Samsung leads in hybrid copper bonding for higher stack counts and improved thermal resistance. SK hynix is collaborating with TSMC to optimize HBM and advanced packaging integration, highlighting the critical role of foundry partnerships.

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 | DECISION LOCUS: BUSINESS / RULES

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

P09

HBM4E: Samsung's 4nm base die and thermal specs for NVIDIA

A01 | Test & inspection | Pilot

Date not stated

FACT

Samsung Electronics has developed an advanced HBM4E thermal management solution featuring a 4nm logic-based die and hybrid copper bonding (HCB), slashing thermal resistance by over 20% compared to HBM3E. This innovation also improves heat dissipation by 30% and energy efficiency by 40%, directly addressing NVIDIA's demand for thinner HBM packages. The company targets customer qualification for 16-layer HBM4E stacks by early 2027, significantly bolstering its competitive edge in the high-performance memory market.

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 19/25 | MATURITY 2/5 | DECISION LOCUS: SYSTEM / ADOPTION

P10

NVHBM: Micron's bandwidth and power specs for NVIDIA

A05 | Volume supply | Product adoption

Date not stated

FACT

Micron Technology announced its custom HBM product, NVHBM, developed for NVIDIA, will utilize outsourced base dies. NVIDIA states NVHBM integrates a customized physical layer (PHY) to reduce I/O PHY package area by up to 67%, enhancing memory bandwidth by up to 30% and cutting HBM power consumption by 15% compared to standard HBM4E. This strategy aims to improve Micron's HBM profitability while strengthening its competitive position in the high-performance AI chip market.

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 19/25 | MATURITY 4/5 | DECISION LOCUS: BUSINESS / RULES

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 01-15

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	HBM4E: Samsung's 4nm base die and thermal specs for NVIDIA	https://eastasiabrief.com/semiconductors/samsung-revamps-hbm-thermal-packaging-challenge-tsmc-supply-lead-559
A02	HBM Packaging: Samsung vs SK Hynix patent competition in 2026	https://finance.biggo.com/news/e9f952bd-5cef-43b9-b4f6-2da0d476dade
A03	zHBM: Samsung's vertical stacking and HBM4 performance specs	https://www.mk.co.kr/en/business/12167424
A04	Intel EMIB: 2.5D packaging for HBM and logic explained	https://www.bestpcbs.com/blog/2026/10/emib/
A05	NVHBM: Micron's bandwidth and power specs for NVIDIA	https://www.thelec.net/news/articleView.html?idxno=14372
A06	Ultra C and Ultra ECP: PLP yield and cost efficiency specs	https://chiplet-marketplace.com/insights/expert-perspectives/overcoming-cleaning-and-plating-challenges-in-panel-level-packaging-and-manufacturing
A07	GlobalFoundries: \$2B TSMC CoWoS interposer US production timeline	https://www.idcnova.com/news/GlobalFoundries-and-TSMC-Sign-2B-AI-Packaging-Deal-to-Scale-CoWoS-Capacity
A08	HBM4: SK hynix and Samsung 16-layer vs NVIDIA 12-layer specs	https://m.koreaherald.com/article/10897751
A09	AMD: Taiwan packaging and substrate capacity for 2027/2028	https://finance.biggo.com/news/8000d2dd-5ad9-4964-96c2-40c848906617
A10	3D Metrology explained: Hybrid bonding defect detection specs	https://semiengineering.com/finding-critical-defects-before-they-become-costly-failures-process-control-for-hybrid-bonding-and-advanced-packaging/
A11	ABF Substrates: 2026 production timeline and layer bottlenecks	https://techsoda.substack.com/p/dr-change-column-15-as-chips-get
A12	Glass Substrates: Intel Clearwater Forest 2026 roadmap	https://www.simpletechtrend.com/en/post/tgv-glass-substrate-2026-advanced-packaging-en
A13	Towa: \$31.6M AI chip packaging factory and 2029 supply specs	http://www.globtechwire.com/articles/towa-plans-new-japan-factory-as-ai-chip-demand-drives-advanced-packaging
A14	Ibiden: 2027 server CPU substrate area and growth projections	https://fundaai.substack.com/p/deepibiden-the-substrate-sellers
A15	Elephantech: Toyosu Packaging Center 2027 launch and specs	https://evertiq.com/design/2026-10-08-elephantech-to-open-advanced-packaging-centre-in-tokyo

Semiconductor_BackEnd — Selected Articles

Date: 2026-10-11

Articles: 15

Table of Contents

- #01 HBM4E: Samsung's 4nm base die and thermal specs for NVIDIA
- #02 HBM Packaging: Samsung vs SK Hynix patent competition in 2026
- #03 zHBM: Samsung's vertical stacking and HBM4 performance specs
- #04 Intel EMIB: 2.5D packaging for HBM and logic explained
- #05 NVHBM: Micron's bandwidth and power specs for NVIDIA
- #06 Ultra C and Ultra ECP: PLP yield and cost efficiency specs
- #07 GlobalFoundries: \$2B TSMC CoWoS interposer US production timeline
- #08 HBM4: SK hynix and Samsung 16-layer vs NVIDIA 12-layer specs
- #09 AMD: Taiwan packaging and substrate capacity for 2027/2028
- #10 3D Metrology explained: Hybrid bonding defect detection specs
- #11 ABF Substrates: 2026 production timeline and layer bottlenecks
- #12 Glass Substrates: Intel Clearwater Forest 2026 roadmap
- #13 Towa: \$31.6M AI chip packaging factory and 2029 supply specs
- #14 Ibiden: 2027 server CPU substrate area and growth projections
- #15 Elephantech: Toyosu Packaging Center 2027 launch and specs

#01 HBM4E: Samsung's 4nm base die and thermal specs for NVIDIA

Published October 06, 2026 East Asia Brief South Korea



OVERVIEW

Samsung Electronics has developed an advanced HBM4E thermal management solution featuring a 4nm logic-based die and hybrid copper bonding (HCB), slashing thermal resistance by over 20% compared to HBM3E. This innovation also improves heat dissipation by 30% and energy efficiency by 40%, directly addressing NVIDIA's demand for thinner HBM packages. The company targets customer qualification for 16-layer HBM4E stacks by early 2027, significantly bolstering its competitive edge in the high-performance memory market.

Key Findings

Samsung Electronics has introduced a groundbreaking thermal management solution for HBM4E, integrating a 4nm logic-based die with hybrid copper bonding (HCB) technology. This development allows Samsung to meet NVIDIA's stringent requirement for a 775-micrometer HBM4 package height, demonstrating a more than 20% reduction in thermal resistance, a 30% improvement in heat dissipation, and a 40% enhancement in energy efficiency compared to HBM3E.

Technical / Clinical Details

The core of Samsung's HBM4E solution lies in its dual technological advancements. Firstly, the use of a cutting-edge 4nm logic process for the base die, which manages data transfer between the HBM and AI accelerators. This not only curtails power consumption and heat generation from the base die but also enables a doubling of I/O terminals compared to HBM3E. Secondly, the implementation of Hybrid Copper Bonding (HCB) fundamentally improves interconnects. Unlike traditional micro-bump-based thermo-compression bonding (TCB), HCB directly bonds copper, achieving ultra-fine pitches below 10 micrometers and significantly boosting thermal conductivity. This results in reduced layer spacing, thinner stacks, and highly efficient heat dissipation, ultimately leading to a 20% plus reduction in overall thermal resistance. Further contributions to energy efficiency and thermal performance come from optimized low-voltage Through-Silicon Vias (TSVs) and an enhanced power delivery network.

Background & Context

As artificial intelligence continues its rapid expansion, HBM has become an indispensable component in high-performance computing, particularly for GPUs and AI accelerators. However, the increasing vertical stacking of HBM layers has exacerbated challenges related to package height and heat generation. Key customers like NVIDIA are pushing for thinner, cooler-running HBM modules, imposing package height standards below the JEDEC specification. Samsung's latest innovation directly addresses these critical market demands, positioning the company to strengthen its competitive standing against rivals like TSMC and SK Hynix in the burgeoning HBM market.

Strategic Significance & Outlook

Samsung plans to begin customer qualification for its 16-layer HBM4E stacks in early 2027, a milestone that could significantly elevate its market presence in next-generation AI semiconductor supply. Beyond HCB, the company is also exploring alternative thermal mitigation techniques such as Heat Path Block (HPB) for HBM4E. This holistic approach to cooling, extending from the chip to the package and even system levels, indicates a strategic vision to push the boundaries of HBM technology. Such advancements are crucial for sustained improvements in computing performance and energy efficiency, vital for the continued growth of the AI era.

Reference technology guide: High Bandwidth Memory (HBM): <https://troy-technical.com/semiconductor-back-end-term-hbm/>

Reference technology guide: Bonding: <https://troy-technical.com/semiconductor-back-end-term-bonding/>

Reference technology guide: Through-Silicon Via (TSV): <https://troy-technical.com/semiconductor-back-end-term-tsv/>

Reference company guide: Samsung Electronics: <https://troy-technical.com/semiconductor-back-end-company-samsung-electronics/>


Reference company guide: NVIDIA: <https://troy-technical.com/semiconductor-back-end-company-nvidia/>

Reference company guide: TSMC: <https://troy-technical.com/semiconductor-back-end-company-tsmc/>

Source: <https://eastasiabrief.com/semiconductors/samsung-revamps-hbm-thermal-packaging-challenge-tsmc-supply-lead-559>

#02 HBM Packaging: Samsung vs SK Hynix patent competition in 2026

Published October 06, 2026 BigGo Finance South Korea



OVERVIEW

Samsung Electronics and SK Hynix have filed competing patents in the U.S. for advanced HBM packaging, aiming to enhance efficiency and reliability. Samsung proposes integrating pre- and post-bonding inspection marks to reduce chip area and improve alignment accuracy, thereby minimizing voids and delamination at bonding interfaces. SK Hynix's patent details a design using polyimide for over 50% of the insulation layer to suppress thermal deformation-induced cracking, reflecting both companies' focus on robust HBM manufacturing.

Key Findings

Samsung Electronics and SK Hynix are at the forefront of the advanced packaging race for High Bandwidth Memory (HBM), each filing innovative patents with the U.S. Patent and Trademark Office. Samsung's proposed technology integrates pre- and post-bonding inspection marks into a single structure, aiming to reduce chip area, significantly enhance alignment precision, and effectively mitigate voids and delamination at the bonding interface. SK Hynix, in parallel, has introduced a design that incorporates polyimide for over 50% of the insulation layer within the HBM stack, effectively suppressing cracks caused by thermal deformation. These patents promise substantial advancements in HBM reliability and manufacturing efficiency.

Technical / Clinical Details

Samsung's patent focuses on streamlining the manufacturing process for stacked chips. By consolidating traditionally separate pre-bonding and post-bonding inspection marks into a singular structure, the technology minimizes the required footprint for these marks, optimizing valuable chip area. This integration also enables highly precise alignment with a single mark, reducing positional inaccuracies during chip stacking. Consequently, it lowers the risk of defects such as voids and delamination at the bonding interfaces, crucial for improving HBM yield rates and ensuring long-term reliability. SK Hynix's innovation, conversely, targets thermal management and structural stability within HBM stacks. Multi-layered HBM dies undergo repeated expansion and contraction during operation due to temperature fluctuations, which can induce stress concentrations and lead to cracking. Their proposed design utilizes polyimide, a high-performance polymer known for its excellent heat resistance and flexibility, for more than 50% of the insulation layer. This material choice effectively mitigates internal stresses caused by thermal deformation, preventing crack formation and significantly enhancing the HBM's durability and reliability for operation in more demanding environments.

Background & Context

The explosive growth of AI and High-Performance Computing (HPC) has made HBM a critical determinant of GPU and AI accelerator performance. As HBM stack counts continue to increase, from HBM3E to HBM4 and beyond, efficient chip area utilization, stacking precision, and robust thermal management have become paramount challenges. The quality of bonding interfaces directly impacts signal integrity and heat dissipation, making defect reduction essential. Furthermore, thermal cycling stress in multi-layered stacks represents a significant concern for package reliability. These patents from Samsung and SK Hynix reflect their strategic efforts to address these challenges and establish technological leadership, thereby accelerating the mass production and adoption of next-generation HBM.

Strategic Significance & Outlook

These advanced packaging patents sought by both companies are expected to contribute not only to HBM manufacturing cost reduction and yield improvement but also to a substantial enhancement in product reliability and durability. This will allow HBM to expand its application scope as a core component of more powerful and stable AI systems. Samsung's integrated inspection mark technology has the potential to become a standard for fine alignment in future chiplet and 3D stacking technologies, while SK Hynix's thermal deformation suppression technology will be crucial for ensuring stable HBM operation in extreme conditions. These technological innovations are poised to solidify both companies' leadership in the AI semiconductor market and positively impact the entire HBM supply chain.

Reference technology guide: High Bandwidth Memory (HBM): <https://troy-technical.com/semiconductor-back-end-term-hbm/>

Reference technology guide: Bonding: <https://troy-technical.com/semiconductor-back-end-term-bonding/>

Reference technology guide: Coefficient of Thermal Expansion (CTE): <https://troy-technical.com/semiconductor-back-end-term-cte/>

Reference company guide: Samsung Electronics: <https://troy-technical.com/semiconductor-back-end-company-samsung-electronics/>


Reference company guide: SK hynix: <https://troy-technical.com/semiconductor-back-end-company-sk-hynix/>

Source: <https://finance.biggo.com/news/e9f952bd-5cef-43b9-b4f6-2da0d476dade>

Collected: October 09, 2026 | Automated Research System (Gemini API)

#03 zHBM: Samsung's vertical stacking and HBM4 performance specs

Published October 02, 2026 MK South Korea



OVERVIEW

Samsung Electronics has doubled HBM4's I/O pin count compared to HBM3E, significantly accelerating data transfer with AI accelerators. The company is developing zHBM, a next-generation vertical stacking technology that places memory directly atop AI accelerators, aiming for up to 8 times higher performance, 3 times greater power efficiency, and over 50% lower thermal resistance than HBM5. This innovative architecture minimizes data transmission distances, potentially revolutionizing AI processing capabilities by overcoming traditional HBM placement limitations.

Key Findings

Samsung Electronics has achieved a significant leap in HBM4 capabilities, effectively doubling the I/O pin count compared to HBM3E, which drastically accelerates data transfer rates between AI accelerators and memory. Furthermore, the company is pioneering a revolutionary next-generation stacked memory technology called 'zHBM,' which vertically stacks memory directly on top of AI accelerators, a departure from traditional HBM configurations positioned alongside. This innovative architecture targets an ambitious eightfold increase in performance, a threefold improvement in power efficiency, and a reduction in thermal resistance to less than half compared to HBM5.

Technical / Clinical Details

Samsung's HBM4 leverages its proprietary 4nm process technology for the base die, enabling the doubling of I/O pins relative to HBM3E. This expansion significantly boosts the data transfer bandwidth between HBM and AI accelerators, crucial for accelerating AI workloads. The most compelling development is zHBM (Z-axis HBM). This technology adopts a paradigm-shifting approach by directly stacking multiple HBM dies vertically onto an AI accelerator chip. Conventional HBM systems typically place HBM modules horizontally adjacent to GPUs or other AI accelerators, connecting them via an interposer. In contrast, zHBM's direct vertical interconnection physically minimizes data transmission distances. This extremely short path dramatically reduces signal latency and power consumption, potentially boosting data transfer speeds and overall system performance by up to eight times. This vertical integration also contributes to a smaller package footprint, aiding in system miniaturization. The target of reducing thermal resistance by more than 50% is anticipated to be achieved through a combination of novel cooling technologies, advanced materials, and optimized structural designs.

Background & Context

In the rapidly expanding fields of AI and High-Performance Computing, there is an ever-increasing demand for greater memory bandwidth and processing speed to handle an explosion of data. While conventional HBM technology brought significant advancements with its stacked structure and wide bandwidth, challenges such as data transmission latency and power consumption—arising from the physical separation between AI accelerators and HBM—became apparent. Specifically, longer data transmission distances increase signal degradation and noise, impacting ultimate system performance. Vertical stacking approaches like zHBM are critical breakthroughs to mitigate this 'Von Neumann bottleneck' and minimize the distance between memory and processor, thereby elevating AI chip performance to the next level. This trend exemplifies the broader semiconductor industry's strategy to surpass Moore's Law limitations and achieve performance gains through advanced packaging and stacking technologies.

Strategic Significance & Outlook


Should Samsung's zHBM technology become commercially viable, it promises to revolutionize AI accelerator design and dramatically enhance AI application capabilities. The projected eightfold performance increase, threefold power efficiency improvement, and significant thermal resistance reduction will have the greatest impact on large-scale AI model training and real-time inference. This will substantially contribute to power consumption reduction and computing power enhancement in data centers, accelerating breakthroughs in diverse AI fields such as autonomous driving, natural language processing, and image recognition. Through zHBM, Samsung aims to establish unparalleled technological leadership in the next-generation AI semiconductor market and foster a new ecosystem.

Reference technology guide: High Bandwidth Memory (HBM): <https://troy-technical.com/semiconductor-back-end-term-hbm/>

Reference company guide: Samsung Electronics: <https://troy-technical.com/semiconductor-back-end-company-samsung-electronics/>

#04 Intel EMIB: 2.5D packaging for HBM and logic explained

Published October 02, 2026 bestpcbs.com USA



OVERVIEW

Intel's EMIB (Embedded Multi-die Interconnect Bridge) is a 2.5D packaging technology that connects adjacent dies via small silicon bridges embedded within the package substrate, achieving high-bandwidth links. While the package substrate handles extensive routing and power delivery, EMIB concentrates fine-pitch routing only where high bandwidth is critical. This localized approach makes EMIB ideal for chiplet processors, AI accelerators, HPC devices, FPGAs, and integrated logic-HBM packages, enabling high-density integration with superior performance.

Key Findings

Intel's Embedded Multi-die Interconnect Bridge (EMIB) is a 2.5D packaging technology that efficiently connects adjacent dies through small silicon bridges embedded within the package substrate. This innovative approach assigns extensive routing and power distribution to the standard package substrate, while reserving ultra-fine pitch routing for critical areas demanding high-bandwidth connectivity. This method eliminates the need for large, expensive silicon interposers across the entire package, thus enabling cost-effective, high-performance multi-die integration.

Technical / Clinical Details

The EMIB architecture is characterized by a compact silicon bridge directly embedded within the dielectric layers of the package substrate. This bridge facilitates ultra-fine pitch wiring between neighboring chiplets, allowing for high-bandwidth, low-latency communication at the chiplet level. Unlike traditional 2.5D packaging solutions that require multiple dies to be placed on a large silicon interposer, EMIB confines this interposer functionality to a localized bridge, significantly reducing manufacturing costs and complexity. For areas outside the bridge, the system relies on a more cost-effective organic package substrate for signal routing and power delivery. This 'localized interposer' strategy offers superior flexibility and scalability by providing high-density connectivity only where necessary, while utilizing standard packaging technologies elsewhere. EMIB is particularly well-suited for chiplet processors, AI accelerators, High-Performance Computing (HPC) devices, FPGAs, and packages combining logic chips with High Bandwidth Memory (HBM), enabling these complex systems to seamlessly collaborate within a single package for superior performance.

Background & Context

As the physical limits of Moore's Law become increasingly apparent, the semiconductor industry has shifted its focus from solely increasing transistor density within a single chip to achieving performance gains through advanced packaging technologies. Chiplet architectures and 2.5D/3D packaging, which integrate multiple functional dies into a single package, have become essential for balancing system performance, power efficiency, and manufacturing costs. Intel's EMIB emerged within this industry trend as a solution that combines high bandwidth with cost efficiency. While competing technologies like TSMC's CoWoS (Chip-on-Wafer-on-Substrate) utilize larger silicon interposers, EMIB's more targeted approach gives it a distinct competitive advantage in specific applications. The tight integration of memory and logic in AI accelerators and HPC is critical for maximizing data processing capabilities, and EMIB provides an effective means to achieve this.

Strategic Significance & Outlook

EMIB technology is expected to play a crucial role in shaping the future of high-performance computing. Particularly with the evolution of AI and the increasing demands for data processing, where closer integration of HBM and logic dies is imperative, technologies like EMIB are key to overcoming bandwidth and power efficiency bottlenecks. In the future, it is anticipated to expand the possibilities for heterogeneous integration, allowing for the combination of more diverse types of chiplets (e.g., analog, RF, security modules) into a single package. This will accelerate the design of custom solutions optimized for specific applications. Intel's continued evolution of EMIB and its promotion of the chiplet ecosystem are expected to drive innovation across the broader semiconductor industry.

Reference technology guide: Package Substrates: <https://troy-technical.com/semiconductor-back-end-term-package-substrate/>

Reference technology guide: Bridges & EMIB: <https://troy-technical.com/semiconductor-back-end-term-bridge/>

Reference technology guide: High Bandwidth Memory (HBM): <https://troy-technical.com/semiconductor-back-end-term-hbm/>

Reference company guide: Intel: <https://troy-technical.com/semiconductor-back-end-company-intel/>


Reference company guide: TSMC: <https://troy-technical.com/semiconductor-back-end-company-tsmc/>

Source: <https://www.bestpcbs.com/blog/2026/10/emib/>

Collected: October 09, 2026 | Automated Research System (Gemini API)

#05 NVHBM: Micron's bandwidth and power specs for NVIDIA

Published October 01, 2026 The Elec Inc. USA



OVERVIEW

Micron Technology announced its custom HBM product, NVHBM, developed for NVIDIA, will utilize outsourced base dies. NVIDIA states NVHBM integrates a customized physical layer (PHY) to reduce I/O PHY package area by up to 67%, enhancing memory bandwidth by up to 30% and cutting HBM power consumption by 15% compared to standard HBM4E. This strategy aims to improve Micron's HBM profitability while strengthening its competitive position in the high-performance AI chip market.

Key Findings

Micron Technology has unveiled a strategic shift for its custom HBM product, "NVHBM," developed specifically for NVIDIA, announcing that it will utilize externally sourced base dies. According to NVIDIA, NVHBM integrates a customized physical layer (PHY), achieving a remarkable reduction of up to 67% in package area required for the I/O PHY. This innovation is projected to boost memory bandwidth by up to 30% and decrease HBM power consumption by 15% compared to standard HBM4E, positioning Micron to enhance HBM product profitability and reinforce its competitiveness in the high-performance AI chip market through this outsourced base die approach.

Technical / Clinical Details

The core innovation of NVHBM lies in its customized physical layer (PHY) integrated into the base die. This PHY circuit, optimized for NVIDIA's specific AI accelerator designs, efficiently handles I/O signal transmission and reception. Such optimization dramatically reduces the I/O PHY area traditionally required in HBM packaging, leading to an impressive up to 67% reduction in overall package area. This reduction directly translates to lower manufacturing costs and increased packaging density. The enhanced efficiency of the I/O PHY also improves signal integrity, leading to higher data transfer rates. Consequently, NVHBM can achieve up to a 30% increase in memory bandwidth compared to standard HBM4E, directly contributing to superior data processing capabilities for AI workloads. Furthermore, the optimized power efficiency, resulting in a 15% reduction in HBM power consumption, holds significant implications for reducing data center operating costs and minimizing environmental impact.

Background & Context

The burgeoning AI semiconductor market has fueled an explosive demand for high-performance HBM. Leading AI chip manufacturers, particularly NVIDIA, require HBM with exceptionally high bandwidth, low power consumption, and compact package sizes to maximize GPU performance. For memory manufacturers, HBM production is complex and costly, making profitability a persistent challenge. Micron's strategy of utilizing externally sourced base dies aims to concentrate its internal resources on memory stack manufacturing while mitigating the manufacturing costs and development risks associated with base dies. This approach allows Micron to supply custom HBM that meets NVIDIA's stringent demands, optimizing cost-efficiency and time-to-market, thereby maintaining and strengthening its competitive edge in the rapidly growing AI market. This trend signifies a shift from vertical integration to horizontal specialization within the semiconductor supply chain, particularly highlighting the increasing importance of advanced packaging and custom solutions.

Strategic Significance & Outlook

Micron's NVHBM strategy is poised to solidify its position in the AI semiconductor market. Collaborating with external foundries for base die procurement enhances HBM production flexibility, enabling rapid adaptation to market demand fluctuations. The projected up to 30% bandwidth increase and 15% power reduction are crucial for boosting the performance of NVIDIA's next-generation AI accelerators, essential for training and inference of large-scale AI models. In the future, the development of such custom HBM products and collaboration with external ecosystems could become a new business model within the memory industry, accelerating HBM evolution in the AI era. This advancement is expected to drive progress in AI technologies across various domains, including data centers, edge AI, and high-performance computing.

Reference technology guide: High Bandwidth Memory (HBM): <https://troy-technical.com/semiconductor-back-end-term-hbm/>

Reference company guide: NVIDIA: <https://troy-technical.com/semiconductor-back-end-company-nvidia/>


Reference company guide: Micron Technology: <https://troy-technical.com/semiconductor-back-end-company-micron/>

Source: <https://www.thelec.net/news/articleView.html?idxno=14372>

Collected: October 09, 2026 | Automated Research System (Gemini API)

#06 Ultra C and Ultra ECP: PLP yield and cost efficiency specs

Published October 01, 2026 chiplet-marketplace.com Global



OVERVIEW

Ultra C vac-p and Ultra ECP ap-p systems have been developed to overcome critical cleaning and plating challenges in Panel-Level Packaging (PLP). Ultra C vac-p is essential for flux residue removal before underfill, eliminating voids even with fine bump pitches below 40µm and large chip sizes. The Ultra ECP ap-p plating system supports multiple applications like pillars, bumps, and RDLs, demonstrating excellent uniformity and adhesion to glass panels. These technologies facilitate the transition to PLP, boosting package count per substrate area, throughput, yield, and significantly reducing costs.

Key Findings

Precision and efficiency in cleaning and plating processes, long-standing hurdles in Panel-Level Packaging (PLP) and manufacturing, have been significantly advanced with the introduction of the Ultra C vac-p cleaning system and Ultra ECP ap-p plating system. Ultra C vac-p effectively removes flux residues prior to underfill, critically eliminating detrimental voids in fine bump pitches below 40µm and large chip sizes. Concurrently, the Ultra ECP ap-p plating system accommodates diverse plating applications such as pillars, bumps, and redistribution layers (RDLs), achieving superior uniformity and robust adhesion to glass panels through its unique multi-anode system and horizontal panel alignment. These innovations are accelerating the shift from circular wafers to large rectangular panels, maximizing package count per substrate area, and dramatically improving throughput, yield, and manufacturing costs.

Technical / Clinical Details

The Ultra C vac-p system leverages a vacuum-based precision cleaning process to reliably remove minute flux residues and other contaminants before underfilling. This is particularly crucial for preventing shorts or voids between fine-pitch bumps in chiplet, 2.5D, and 3D structures. Its efficacy is pronounced for bump pitches below 40µm and large chip sizes, substantially enhancing connection reliability. Meanwhile, the Ultra ECP ap-p plating system employs a unique multi-anode configuration and horizontal panel alignment during its electroplating process. The multi-anode system precisely controls the electric field distribution, ensuring exceptionally high uniformity of the plated film thickness across the entire panel. The horizontal panel arrangement optimizes plating solution flow, mitigating the risk of air bubble entrapment and particle adhesion. This system supports various advanced packaging plating processes, including pillar, micro-bump, and RDL formation, and demonstrates excellent adhesion, especially to novel substrate materials like glass panels. These technologies enable the manufacturing of high-quality, reliable packages in PLP with increased throughput and reduced costs.

Background & Context

As semiconductor packaging technologies confront the physical limits of Moore's Law, they have become a primary driver for enhancing system performance and reducing costs. The escalating demand for high-performance applications such as AI, HPC, and mobile devices necessitates denser, smaller, and more power-efficient packaging solutions. Panel-Level Packaging (PLP) offers significant advantages over traditional circular wafer-based packaging by utilizing large, rectangular panels, which improves material utilization efficiency and substantially reduces manufacturing costs. However, the adoption of PLP has been met with technical challenges, including achieving uniform cleaning across large substrates, precise plating, and compatibility with new materials (e.g., glass substrates). Advanced process solutions like Ultra C vac-p and Ultra ECP ap-p are indispensable for overcoming these hurdles and enabling the widespread adoption of PLP.

Strategic Significance & Outlook

The emergence of innovative technologies such as Ultra C vac-p and Ultra ECP ap-p significantly increases the likelihood of Panel-Level Packaging becoming a standard for next-generation semiconductor manufacturing. PLP's efficiency and cost advantages will accelerate the proliferation of advanced packaging technologies like chiplets, 2.5D/3D integration, and Fan-Out Panel-Level Packaging (FOPLP). This is expected to reduce manufacturing costs for high-performance chips such as AI accelerators and HBM, promoting their wider application. Furthermore, the capability to accommodate new materials like glass substrates lays a crucial foundation for future packaging technologies demanding ultra-fine wiring and high-speed signal transmission. These advancements will play a vital role in enhancing manufacturing efficiency and competitiveness across the entire semiconductor industry supply chain.

Reference technology guide: Panel-Level Packaging (PLP): <https://troy-technical.com/semiconductor-back-end-term-plp/>

Reference technology guide: Bumps: <https://troy-technical.com/semiconductor-back-end-term-bump/>


Reference technology guide: Underfill & Encapsulants: <https://troy-technical.com/semiconductor-back-end-term-underfill/>

Source: <https://chiplet-marketplace.com/insights/expert-perspectives/overcoming-cleaning-and-plating-challenges-in-panel-level-packaging-and-manufacturing>

Collected: October 09, 2026 | Automated Research System (Gemini API)

#07 GlobalFoundries: \$2B TSMC CoWoS interposer US production timeline

Published October 09, 2026 IDC Nova USA



OVERVIEW

GlobalFoundries and TSMC have signed a 5-year, \$2 billion agreement for GlobalFoundries to manufacture silicon interposers for TSMC's CoWoS advanced packaging ecosystem. This strategic partnership will significantly expand GlobalFoundries' production capacity at its Malta, New York facility, with mass production slated to begin in early 2028. The deal addresses a critical bottleneck in the AI semiconductor supply chain by establishing a domestic US source for these essential components.

Key Findings

GlobalFoundries and the world's leading foundry, TSMC, have entered into a five-year, \$2 billion agreement to scale the production of silicon interposers for TSMC's CoWoS (Chip-on-Wafer-on-Substrate) advanced packaging technology. GlobalFoundries will expand its existing facility in Malta, New York, with mass production anticipated to commence in the first half of 2028, significantly bolstering the supply chain for critical AI components.

Technical / Clinical Details

Silicon interposers are indispensable for high-density integration of HBM (High Bandwidth Memory) and high-performance logic dies within AI GPUs and accelerators. CoWoS technology leverages these interposers to stack multiple HBM modules and processors, enabling unparalleled data bandwidth and power efficiency crucial for advanced AI workloads. Historically, the complex and time-consuming manufacturing of interposers has been a primary constraint on AI chip supply. GlobalFoundries' increased capacity aims to alleviate this bottleneck and enhance supply chain resilience, particularly within the United States.

Background & Context

The burgeoning demand for AI has led to an unprecedented surge in high-performance AI chip requirements, severely straining advanced packaging and HBM memory production capacities. TSMC's CoWoS lines are currently booked through 2027 by major AI chipmakers like NVIDIA, with lead times extending to 52-78 weeks. This agreement aligns with the US CHIPS Act's objectives to promote domestic semiconductor manufacturing, establishing a secure, in-country supply for a vital component and mitigating geopolitical supply risks.

Strategic Significance & Outlook

GlobalFoundries' mass production of silicon interposers in the US will diversify the global AI semiconductor supply chain and accelerate AI technology development within the country. This collaboration is expected to not only contribute to the stable supply of AI chips but also mark a significant step for GlobalFoundries in enhancing its competitiveness in the advanced packaging market. As AI demand continues to soar, securing packaging technology and manufacturing capacity will remain a pivotal challenge influencing the growth of the entire semiconductor industry.

Reference technology guide: CoWoS: <https://troy-technical.com/semiconductor-back-end-term-cowos/>

Reference technology guide: High Bandwidth Memory (HBM): <https://troy-technical.com/semiconductor-back-end-term-hbm/>

Reference company guide: TSMC: <https://troy-technical.com/semiconductor-back-end-company-tsmc/>

Reference company guide: NVIDIA: <https://troy-technical.com/semiconductor-back-end-company-nvidia/>

Source: <https://www.idcnova.com/news/GlobalFoundries-and-TSMC-Sign-2B-AI-Packaging-Deal-to-Scale-CoWoS-Capacity>

#08 HBM4: SK hynix and Samsung 16-layer vs NVIDIA 12-layer specs

Published October 09, 2026 The Korea Herald South Korea



OVERVIEW

Despite NVIDIA's Vera Rubin AI platform favoring 12-layer HBM4 chips, SK hynix and Samsung Electronics are aggressively pursuing 16-layer HBM4 technology. Samsung leads in hybrid copper bonding for higher stack counts and improved thermal resistance. SK hynix is collaborating with TSMC to optimize HBM and advanced packaging integration, highlighting the critical role of foundry partnerships. This pursuit is fueled by AI demand consuming over 3x more wafer capacity for HBM than conventional DRAM, portending continued tight DRAM supply.

Key Findings

Even as NVIDIA's forthcoming Vera Rubin AI platform is expected to rely predominantly on 12-layer HBM4 memory chips, South Korean memory giants SK hynix and Samsung Electronics are pushing ahead with the development of more advanced 16-layer HBM4 technology. Samsung, in particular, has made strides with its hybrid copper bonding technique, which promises to support memory stacks beyond 16 layers and significantly reduce thermal resistance, positioning it ahead in this crucial innovation race.

Technical / Clinical Details

High Bandwidth Memory (HBM) involves stacking multiple DRAM dies vertically, connected via an interposer to a CPU or GPU, offering substantially higher bandwidth than traditional DRAM. The increasing demand for greater data throughput and capacity in AI applications necessitates higher stack counts. While NVIDIA's Rubin platform targets 12 layers, SK hynix and Samsung are developing 16-layer HBM4 to meet future AI chip requirements. Samsung's hybrid copper bonding technology, a key innovation, eliminates traditional chip bumps, directly bonding copper interconnects between silicon dies. This results in shorter connection distances, lower latency, enhanced performance, reduced signal distortion, and crucially, enables the physical feasibility of stacking 16 or more layers. SK hynix is optimizing HBM product integration with advanced packaging by collaborating closely with TSMC, emphasizing that HBM performance is not solely dependent on the memory itself but also on foundry packaging expertise.

Background & Context

The exponential growth in AI demand has fundamentally reshaped the DRAM industry, with HBM becoming a pivotal component. HBM manufacturing is considerably more complex than standard DRAM, yielding significantly fewer memory bits per wafer. High-performance AI chip production consumes over three times the wafer capacity compared to conventional DRAM. This disparity suggests that continued robust AI demand will likely sustain tight DRAM supply conditions across the industry. Major AI chip designers are therefore intensifying collaborations with memory manufacturers and foundries to secure stable supplies of high-performance HBM.

Strategic Significance & Outlook

The fierce competition between SK hynix and Samsung in 16-layer HBM4 development underscores the critical race for AI semiconductor performance leadership and market share. Innovative packaging technologies like hybrid copper bonding are indispensable for achieving higher HBM capacity, superior performance, and improved power efficiency, forming the bedrock of next-generation AI infrastructure. Bridging the gap between major customer roadmaps, such as NVIDIA's, and memory manufacturers' technological advancements will be crucial for market dynamics. As HBM's proportion within the overall DRAM market increases, it is poised to further alter supply structures and significantly impact pricing strategies.

Reference technology guide: High Bandwidth Memory (HBM): <https://troy-technical.com/semiconductor-back-end-term-hbm/>

Reference technology guide: Bonding: <https://troy-technical.com/semiconductor-back-end-term-bonding/>

Reference technology guide: Bumps: <https://troy-technical.com/semiconductor-back-end-term-bump/>

Reference company guide: Samsung Electronics: <https://troy-technical.com/semiconductor-back-end-company-samsung-electronics/>


Reference company guide: SK hynix: <https://troy-technical.com/semiconductor-back-end-company-sk-hynix/>

Reference company guide: NVIDIA: <https://troy-technical.com/semiconductor-back-end-company-nvidia/>

Source: <https://m.koreaherald.com/article/10897751>

#09 AMD: Taiwan packaging and substrate capacity for 2027/2028

Published October 06, 2026 BigGo Finance Taiwan



OVERVIEW

AMD CEO Lisa Su visited Taiwan to secure advanced process technology, packaging, and substrate capacity for 2027 and 2028, ensuring smooth shipments of next-generation AI server products. While TSMC accelerates 2nm and CoWoS-L expansion, the actual bottlenecks remain HBM memory supply, bit output degradation, and silicon interposer/package test yields. AMD's MI450 series AI platform, utilizing TSMC's 2nm and CoWoS, is slated for higher volume shipments in 2027.

Key Findings

AMD CEO Lisa Su embarked on a rapid visit to Taiwan to lock in critical production capacity for advanced processes, packaging, and substrates, with a strategic focus on 2027 and 2028. This move aims to ensure the smooth rollout of AMD's next-generation AI server products, highlighting the immense pressure AI chip demand is placing on the entire semiconductor supply chain.

Technical / Clinical Details

AMD's MI450 series AI platform leverages TSMC's cutting-edge 2nm process technology and CoWoS (Chip-on-Wafer-on-Substrate) packaging, which is essential for high-performance AI accelerators. Advanced packaging techniques like CoWoS-L integrate logic dies with multiple HBM (High Bandwidth Memory) stacks on a silicon interposer, achieving extremely high data bandwidth and power efficiency. Although TSMC has announced accelerated expansion of its 2nm and CoWoS-L capacities, the actual supply bottlenecks persist in areas such as HBM memory availability (despite long-term contracts), degraded bit output, and yield challenges in silicon interposers and packaging tests. Securing HBM memory, in particular, directly dictates the final supply volume of AI accelerators.

Background & Context

The explosive growth of AI demand has placed unprecedented strain on the semiconductor industry. Specifically, advanced process nodes, sophisticated packaging, and HBM memory have emerged as the 'new bottlenecks' in AI chip manufacturing. Demand for CoWoS packaging from major AI chipmakers like NVIDIA reportedly consumes a significant portion of TSMC's capacity, with lead times stretching to 52-78 weeks. AMD, striving to expand its market share, faces the urgent challenge of securing production capacity in this highly constrained supply environment. CEO Su's visit underscores how top-tier AI chip design companies prioritize supply chain stability.

Strategic Significance & Outlook

With AMD planning even higher volume shipments of its MI450 series AI platform in 2027, securing advanced process and packaging capacity is paramount for its growth strategy. If bottlenecks such as HBM memory, interposers, and packaging test yields are not resolved, the overall growth of the AI chip market could be hampered. Close collaboration between AMD and TSMC is key to overcoming these challenges and meeting the demand for next-generation AI infrastructure. Such strategic alliances and investments will continue to be crucial movements toward stabilizing the AI semiconductor supply chain.

Reference technology guide: CoWoS: <https://troy-technical.com/semiconductor-back-end-term-cowos/>

Reference technology guide: High Bandwidth Memory (HBM): <https://troy-technical.com/semiconductor-back-end-term-hbm/>

Reference technology guide: Semiconductor Test: <https://troy-technical.com/semiconductor-back-end-term-test/>

Reference company guide: AMD: <https://troy-technical.com/semiconductor-back-end-company-amd/>

Reference company guide: TSMC: <https://troy-technical.com/semiconductor-back-end-company-tsmc/>


Reference company guide: NVIDIA: <https://troy-technical.com/semiconductor-back-end-company-nvidia/>

Source: <https://finance.biggo.com/news/8000d2dd-5ad9-4964-96c2-40c848906617>

Collected: October 09, 2026 | Automated Research System (Gemini API)

#10 3D Metrology explained: Hybrid bonding defect detection specs

Published October 08, 2026 Semiconductor Engineering USA



OVERVIEW

3D metrology is critically enhancing defect detection in advanced semiconductor packaging, particularly for hybrid bonding, to prevent costly failures. This technology precisely measures fine bump heights and surface profiles in complex structures like chiplets and HBM stacks. By identifying minute defects early in the assembly process, it significantly reduces risks to larger assemblies and associated material/process waste, crucial for improving yields and reducing manufacturing costs in next-generation semiconductors.

Key Findings

In advanced packaging technology, 3D metrology is significantly improving defect detection capabilities, proactively preventing costly failures in hybrid bonding and complex chiplet integrations. This technology enables high-precision measurement of minute bump heights and surface profiles, facilitating early identification of small defects that might otherwise go unnoticed by conventional inspection methods.

Technical & Clinical Details

The assembly of chiplets, high-bandwidth memory (HBM) stacks, and hybrid-bonded devices in advanced semiconductor packaging involves substantial material and process resources. Undetected defects at early stages of this intricate manufacturing flow can jeopardize the entire larger assembly, leading to significant financial losses. For instance, micron-level variations in fine bump heights or nanometer-scale surface roughness on bonding interfaces can cause performance degradation and reliability issues. 3D metrology non-destructively and rapidly scans these microstructures, acquiring three-dimensional shape data. This allows for precise localization, sizing, and characterization of defects. This capability enables manufacturers to more thoroughly characterize increasingly complex structures and enhance process control.

Background & Context

As Moore's Law approaches its physical limits, the semiconductor industry is shifting towards advanced packaging technologies to drive performance improvements and cost reductions. However, techniques like multi-layer chiplet stacking and hybrid bonding dramatically increase manufacturing complexity and, concurrently, the difficulty of defect detection. The impact is particularly severe when defects are discovered after expensive materials and lengthy processes have been invested. Therefore, the ability to accurately detect defects early in the assembly process is a critical factor for ensuring the yield and reliability of advanced packaging.

Strategic Significance & Outlook

The evolution of 3D metrology is key to overcoming challenges faced by semiconductor manufacturers, enabling the mass production of next-generation high-performance chips. Future advancements are expected in faster and more precise 3D metrology equipment, alongside integration with AI-powered data analytics for real-time process control. This will further boost manufacturing yields and improve overall semiconductor quality and cost efficiency. Ultimately, this will accelerate the development of fields reliant on advanced semiconductors, such as AI and High-Performance Computing (HPC).

Reference technology guide: Bonding: <https://troy-technical.com/semiconductor-back-end-term-bonding/>


Reference technology guide: Chiplets: <https://troy-technical.com/semiconductor-back-end-term-chiplet/>

Reference technology guide: Bumps: <https://troy-technical.com/semiconductor-back-end-term-bump/>

Source: <https://semiengineering.com/finding-critical-defects-before-they-become-costly-failures-process-control-for-hybrid-bonding-and-advanced-packaging/>

#11 ABF Substrates: 2026 production timeline and layer bottlenecks

Published October 07, 2026 TechSoda USA



OVERVIEW

As AI accelerators grow in size, ABF (Ajinomoto Build-up Film) substrates are emerging as a critical supply bottleneck in semiconductor back-end processing. Next-generation AI chips are expected to utilize ABF substrates with 18-20 layers by 2026, significantly increasing manufacturing steps like lamination, laser drilling, electroplating, etching, and alignment, thereby extending overall production time. This extended production cycle, coupled with fewer substrates per panel, is becoming a key determinant of supply-demand for advanced AI chips, making efficient ABF substrate capacity crucial for future AI chip market growth.

Key Findings

The evolution and increasing size of AI accelerators have led to ABF (Ajinomoto Build-up Film) substrates becoming a critical bottleneck in the semiconductor supply chain. Specifically, by 2026, next-generation AI chips are projected to require ABF substrates with 18 to 20 layers, which will inevitably lead to significantly extended production times and reduced effective manufacturing capacity.

Technical & Clinical Details

AI accelerators integrate an increasing number of computing dies and High-Bandwidth Memory (HBM), leading to a continuous expansion in package size. Concurrently, the area of the supporting ABF substrate also grows, reducing the number of substrates that can be placed on a single production panel. The multi-layering of ABF substrates necessitates a greater number of complex steps for each additional layer, including lamination, laser drilling, electroplating, etching, and alignment. For instance, while current common ABF substrates typically have 12-16 layers, there is a projection for this to increase to 18-20 layers by 2026. This increase in layer count exponentially extends production time.

Background & Context

The rapid advancement of AI has generated unprecedented demand for high-performance semiconductor chips. However, performance improvements are hitting the limits of traditional miniaturization, increasing reliance on advanced packaging, especially large, multi-layer ABF substrates. The supply capacity for ABF substrates is limited, and while manufacturers are striving to expand capacity to meet growing demand, technical complexity and long manufacturing cycles pose significant challenges. This substrate shortage directly translates into supply constraints for AI chips, potentially decelerating overall market growth.

Strategic Significance & Outlook

The effective production capacity that manufacturers can achieve for large, high-layer-count ABF substrates for advanced AI chips will be a crucial indicator determining the future supply-demand balance of AI chips. To alleviate this bottleneck, innovation in ABF substrate manufacturing processes, expansion of production facilities, and the introduction of new material technologies are urgently required. Furthermore, advancements in alternative technologies, such as glass substrates, are also gaining attention as potential long-term solutions to this challenge.

Reference technology guide: Build-Up Film: <https://troy-technical.com/semiconductor-back-end-term-buildup-film/>

Reference technology guide: High Bandwidth Memory (HBM): <https://troy-technical.com/semiconductor-back-end-term-hbm/>

Reference technology guide: Glass Substrates: <https://troy-technical.com/semiconductor-back-end-term-glass-substrate/>


Reference company guide: Ajinomoto: <https://troy-technical.com/semiconductor-back-end-company-ajinomoto/>

Source: <https://techsoda.substack.com/p/dr-change-column-15-as-chips-get>

Collected: October 09, 2026 | Automated Research System (Gemini API)

#12 Glass Substrates: Intel Clearwater Forest 2026 roadmap

Published October 06, 2026 vertexaisearch.cloud.google.com USA



OVERVIEW

In 2026, facing the physical limits of organic substrates, the semiconductor industry is widely adopting glass substrates for advanced packaging. Intel plans to ship its "Clearwater Forest" processor, the first with a glass core substrate, while TSMC will launch its CoWoS glass pilot line mid-year, and Samsung and Absolics are commencing new factory constructions. Glass is expected to underpin performance and reliability for large packages due to its silicon-matched CTE, superior flatness, and extremely low dielectric loss.

Key Findings

The semiconductor industry is accelerating its adoption of glass substrates as a critical material for advanced packaging, driven by the physical limitations of organic substrates encountered in 2026. Intel is set to mass-produce and ship its "Clearwater Forest" processor, which will be the first to utilize a glass-core substrate. Concurrently, TSMC is scheduled to launch its CoWoS glass pilot line by mid-year, and both Samsung and Absolics have initiated construction on new glass substrate factories, signifying a broad industry shift towards the commercialization of glass substrates.

Technical & Clinical Details

Organic substrates are nearing their physical and electrical limits in meeting the demands for larger, highly integrated packages required by high-performance computing (HPC) and AI applications. In contrast, glass substrates offer several decisive advantages. Firstly, glass has a coefficient of thermal expansion (CTE) that closely matches silicon, significantly reducing the risk of reliability degradation caused by stress between dissimilar materials. Secondly, its inherently superior flatness enables finer pattern formation and higher-density interconnections. Furthermore, glass exhibits extremely low dielectric loss and electrical resistance, which improves power efficiency and signal integrity in high-speed data transmission. These characteristics, particularly for large-area packages, facilitate Through-Glass Via (TGV) technology, forming the foundation for achieving smaller, higher-performance stacked structures.

Background & Context

The escalating demand for AI and high-performance computing has driven an unprecedented increase in semiconductor chip complexity and package size. This has given rise to electrical and thermal challenges that are difficult for conventional organic substrates to address. The introduction of glass substrates is positioned as a critical breakthrough to solve these challenges, enabling the realization of next-generation AI processors with more integrated chiplets and HBM. The successive investments and production line launches by major industry players indicate a significant turning point, moving glass substrates from mere R&D to full-scale commercial production.

Strategic Significance & Outlook

The market introduction of Intel's "Clearwater Forest" processor will serve as a crucial test for the reliability and practicality of glass-core substrate technology. The entry of major foundries like TSMC and Samsung will establish the glass substrate supply chain and bolster mass production capabilities. This trajectory suggests that glass substrates will become indispensable for future AI, data center, and high-performance computing sectors, potentially setting a new standard for semiconductor packaging technology. Over the next few years, advanced packaging solutions utilizing glass substrates are expected to rapidly proliferate, driving industry-wide technological innovation.

Reference technology guide: Glass Substrates: <https://troy-technical.com/semiconductor-back-end-term-glass-substrate/>

Reference technology guide: CoWoS: <https://troy-technical.com/semiconductor-back-end-term-cowos/>

Reference technology guide: Coefficient of Thermal Expansion (CTE): <https://troy-technical.com/semiconductor-back-end-term-cte/>

Reference company guide: Intel: <https://troy-technical.com/semiconductor-back-end-company-intel/>

Reference company guide: TSMC: <https://troy-technical.com/semiconductor-back-end-company-tsmc/>

Reference company guide: Samsung Electronics: <https://troy-technical.com/semiconductor-back-end-company-samsung-electronics/>

Source: <https://www.simpletechtrend.com/en/post/tgv-glass-substrate-2026-advanced-packaging-en>

#13 Towa: \$31.6M AI chip packaging factory and 2029 supply specs

Published October 07, 2026 Globtechwire Japan



OVERVIEW

Japanese semiconductor equipment manufacturer Towa announced plans to invest approximately ¥5 billion (\$31.6 million) in a new factory in Kyoto Prefecture to meet increasing demand for advanced semiconductor packaging equipment driven by AI infrastructure expansion. Aiming for operations to begin in 2029, this new facility will produce high-precision equipment used in the assembly and packaging processes of advanced semiconductors, including those for AI applications, thereby strengthening Japan's contribution to the global supply chain.

Key Findings

Towa, a leading Japanese semiconductor equipment manufacturer, has announced plans to invest approximately ¥5 billion (\$31.6 million) in a new factory in Kyoto Prefecture. This investment aims to meet the escalating demand for advanced semiconductor packaging equipment, fueled by the rapid expansion of AI infrastructure. Expected to commence operations in 2029, this factory will significantly enhance the supply capacity for advanced packaging equipment critical for AI application chips.

Technical & Clinical Details

AI semiconductors require extremely sophisticated packaging technologies to unleash their full performance potential. This includes complex processes such as CoWoS (Chip on Wafer on Substrate) integration of multiple dies and hybrid bonding. Towa's new factory will produce high-precision equipment used in these advanced packaging processes, such as die bonders and molding equipment. The performance of AI chips is intimately linked to packaging quality, making high-precision equipment indispensable for improving yield and ensuring reliability. The new factory will incorporate the latest manufacturing technologies and automation to maximize production efficiency and quality, thereby meeting growing global demand.

Background & Context

The evolution of AI technology is bringing about significant transformations in the semiconductor industry, with a particularly explosive increase in demand for high-performance AI chips. Concurrently, the importance of back-end packaging processes has grown unprecedentedly, raising concerns that the supply capacity of advanced packaging equipment could become a bottleneck for the entire semiconductor supply chain. Towa's current investment underscores the significant role played by the Japanese semiconductor equipment industry in the global AI chip boom. It holds strategic significance by supporting the development of global AI technology through enhanced supply capabilities.

Strategic Significance & Outlook

With the new factory slated for operation in 2029, Towa is expected to further solidify its leading position in the AI chip packaging market. This investment will contribute to the stable supply of advanced packaging equipment, supporting the expansion of AI chip production and thereby promoting the further proliferation and evolution of AI technology. Furthermore, it is anticipated to contribute to enhancing the resilience of the domestic semiconductor supply chain, aligning with the Japanese government's initiatives to strengthen the semiconductor industry. This new factory will symbolize Japan's manufacturing competitiveness in the age of AI.

Reference technology guide: Bonding: <https://troy-technical.com/semiconductor-back-end-term-bonding/>

Reference technology guide: CoWoS: <https://troy-technical.com/semiconductor-back-end-term-cowos/>

Reference technology guide: Package Substrates: <https://troy-technical.com/semiconductor-back-end-term-package-substrate/>

Source: <http://www.globtechwire.com/articles/towa-plans-new-japan-factory-as-ai-chip-demand-drives-advanced-packaging>

#14 Ibiden: 2027 server CPU substrate area and growth projections

Published October 05, 2026 FUNDA | Substack USA



OVERVIEW

Ibiden is leveraging its market position due to a severe shortage of qualified high-end ABF substrates, which cannot keep pace with the increasing area and layer counts of AI silicon. Server CPU shipments from Intel and AMD are projected to increase by 36% and 62% respectively in 2027, driven by new platforms requiring 56%-63% larger substrate areas, further widening the ABF substrate gap. High-end GPUs and EMIB-T designs are consuming premium capacity, with specific architectures like Rubin yielding only 12 units per panel and TPU V9/Trainium 4 only 4 units.

Key Findings

Ibiden is significantly reinforcing its market dominance, driven by a critical shortage of qualified high-end Ajinomoto Build-up Film (ABF) substrates, which are failing to keep pace with the increasing area and layer counts demanded by AI silicon. This scarcity has emerged as a severe bottleneck in the production of AI-specific semiconductors, particularly high-performance GPUs and CPUs, leading to a surge in demand for Ibiden's products.

Technical / Clinical Details

As AI technologies evolve, advanced AI chip packages, integrating chiplets and High Bandwidth Memory (HBM), require increasingly larger and multi-layered ABF substrates. Projections indicate that Intel's and AMD's server CPU shipments will see substantial increases—36% and 62% respectively—by 2027. This growth is driven by new platforms that will demand 56%-63% larger substrate areas compared to existing models, exacerbating the ABF substrate supply gap for CPUs. Furthermore, high-end GPUs are consuming a disproportionate share of premium ABF capacity. For instance, NVIDIA's Rubin architecture is reported to yield only 12 units per panel, while Google's TPU V9 and Trainium 4 achieve only 4 units. This highlights the extreme complexity and extended manufacturing cycles required for these high-performance chips.

Background & Context

ABF substrates are indispensable materials for establishing electrical connections between dies and Printed Circuit Boards (PCBs) in complex semiconductor packages for high-performance CPUs and GPUs. The exponential growth of AI has led to a surge in demand for advanced packaging solutions like CoWoS, which integrate HBM and chiplets, thereby creating an unprecedented demand for supporting ABF substrates. Manufacturing high-end ABF substrates, which require stringent yield and reliability, involves highly sophisticated technology and lengthy qualification processes, intensifying the supply shortage.

Strategic Significance & Outlook

The shortage of ABF substrates is anticipated to persist as a major bottleneck, constraining the expansion of AI chip production capacity in the short term. Key suppliers like Ibiden are expected to maintain strong pricing power and profitability in this market environment. Moving forward, while suppliers will likely invest in capacity expansion, AI chip manufacturers may pursue strategies such as qualifying multiple suppliers, optimizing substrate designs, or exploring alternative materials to secure stable supplies. This supply-demand imbalance will continue to have a significant impact across the entire semiconductor supply chain, dictating the pace of AI hardware innovation and deployment.

Reference technology guide: Chiplets: <https://troy-technical.com/semiconductor-back-end-term-chiplet/>

Reference technology guide: High Bandwidth Memory (HBM): <https://troy-technical.com/semiconductor-back-end-term-hbm/>

Reference technology guide: Build-Up Film: <https://troy-technical.com/semiconductor-back-end-term-buildup-film/>

Reference company guide: Ibiden: <https://troy-technical.com/semiconductor-back-end-company-ibiden/>

Reference company guide: Intel: <https://troy-technical.com/semiconductor-back-end-company-intel/>

Reference company guide: AMD: <https://troy-technical.com/semiconductor-back-end-company-amd/>

Source: <https://fundaai.substack.com/p/deepibiden-the-substrate-sellers>

#15 Elephantech: Toyosu Packaging Center 2027 launch and specs

Published October 08, 2026 Evertiq Japan



OVERVIEW

Japanese PCB and substrate technology developer Elephantech announced it will open its new Toyosu Advanced Packaging Center in Tokyo by March 2027. This center aims to significantly scale up manufacturing and R&D for high-performance PCBs and package substrates specifically designed for AI servers. Leveraging proprietary nanomaterials and precision inkjet processes, Elephantech seeks to meet the escalating demand for high-density interconnects required by next-generation AI infrastructure.

Key Findings

Elephantech, a Japanese developer of PCB and substrate technologies, is set to open its new Toyosu Advanced Packaging Center in Tokyo by March 2027. This strategic investment aims to substantially scale up manufacturing and bolster R&D capabilities for PCBs and package substrates critical to AI servers. The move positions Elephantech to address the growing bottleneck in advanced packaging materials, which is crucial for maximizing AI chip performance.

Technical Details

The new center will utilize Elephantech's proprietary nanomaterials and precision inkjet processes, enabling the production of PCBs and package substrates with significantly higher interconnect density than conventional methods. This technological edge is crucial for meeting the extreme data transfer rates and power efficiency demands of modern AI chips and accelerators. The advanced processes allow for finer line widths and spaces, which are essential for integrating complex AI processors with high-bandwidth memory (HBM). This innovation directly tackles the physical limitations of current packaging technologies, enabling more compact and efficient AI hardware.

Background & Context

The global demand for AI servers and high-speed networking has surged since late 2025, leading to widespread supply constraints for high-performance PCBs, IC substrates, and key raw materials. Reports indicate limited signs of supply improvement even into 2026, with persistent shortages in advanced materials like ABF substrates, process equipment, and consumables. Elephantech's investment in the Toyosu center directly responds to this critical supply chain challenge, aiming to strengthen Japan's position in the global semiconductor ecosystem. The Japanese Ministry of Economy, Trade and Industry (METI) has also committed over 10 trillion yen in public funding by fiscal year 2030 to fortify AI and semiconductor industrial infrastructure, aligning with Elephantech's strategic expansion.

Strategic Significance & Outlook

The operationalization of the Toyosu Advanced Packaging Center is expected to significantly enhance Elephantech's competitive edge in the AI infrastructure market, expanding its supply capacity to both domestic and international clients. Beyond manufacturing, the center will serve as a vital R&D hub, driving the development of next-generation packaging technologies required by future, even more powerful AI chips. This initiative underscores Japan's commitment to maintaining its technological leadership in the semiconductor equipment and materials sector, ensuring its continued indispensable role in the global AI boom.

Reference technology guide: Package Substrates: <https://troy-technical.com/semiconductor-back-end-term-package-substrate/>

Reference technology guide: High Bandwidth Memory (HBM): <https://troy-technical.com/semiconductor-back-end-term-hbm/>

Reference technology guide: Build-Up Film: <https://troy-technical.com/semiconductor-back-end-term-buildup-film/>

Reference company guide: Ajinomoto: <https://troy-technical.com/semiconductor-back-end-company-ajinomoto/>

Source: <https://evertiq.com/design/2026-10-08-elephantech-to-open-advanced-packaging-centre-in-tokyo>