

TROY TECHNICAL

SEMICONDUCTOR BACK-END WEEKLY REPORT

October 17 - October 23, 2026 | Issue 2026-W42

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


12

ANALYZED ARTICLES

10

PRIORITY THEMES

Issued 2026-10-23 | English edition | fixed approved cover artwork



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

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

SIGNAL 01 | MATERIALS

NVIDIA Supply Chain: Capacity Expansion Roadmap for 2027-2028

NVIDIA is partnering with key suppliers to aggressively expand manufacturing capacity for substrates, advanced packaging, testing, and optical engines to meet AI infrastructure demand.

SIGNAL 02 | BONDING & ASSEMBLY

Intel 300mm InP Silicon Photonics: CPO Production Roadmap 2026

Intel claims to be the only high-volume foundry capable of supplying 300mm indium phosphide (InP) silicon photonics at scale.

SIGNAL 03 | ADVANCED PACKAGE

Photonics-Electronics Convergence: NEDO Project for AI Chiplets

NTT Innovative Devices leads a four-company consortium selected by NEDO to develop mass-production technologies for photonics-electronics convergence (PEC) devices.

SIGNAL 04 | TEST & INSPECTION

Samsung: TSMC 3nm process for Nvidia NVHBM base dies

Samsung Electronics will utilize TSMC's 3-nanometre process to produce base dies for Nvidia's next-generation custom high-bandwidth memory (HBM) chips.

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

2.5x A06	SOURCE-BOUND CONDITION Key targets include a 2.5x increase in substrate capacity over three years and a 300% increase in advanced packaging capacity by 2028.
300mm A03	SOURCE-BOUND CONDITION Intel claims to be the only high-volume foundry capable of supplying 300mm indium phosphide (InP) silicon photonics at scale.
30% A10	SOURCE-BOUND CONDITION The initiative targets a reduction in thermal resistance of more than 30% to improve heat dissipation in next-generation memory.
30 month S A07	SOURCE-BOUND CONDITION Unimicron warns of a critical IC substrate crunch characterized by equipment lead times stretching to 30 months and a projected 20% material supply gap by 2027.
35% A01	SOURCE-BOUND CONDITION Google is maintaining its commitment to Intel's EMIB-T packaging for the HumuFish TPU v9 variant despite reported substrate yields of 35% in 3Q26.
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 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

NVIDIA Supply Chain: Capacity Expansion Roadmap for 2027-2028

A06 | Materials | Product adoption

Date not stated

FACT

NVIDIA is partnering with key suppliers to aggressively expand manufacturing capacity for substrates, advanced packaging, testing, and optical engines to meet AI infrastructure demand. Key targets include a 2.5x increase in substrate capacity over three years and a 300% increase in advanced packaging capacity by 2028. These investments aim to resolve capacity constraints and strengthen global supply chain resiliency for AI supercomputing systems.

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 23/25 | MATURITY 4/5 | DECISION LOCUS: BUSINESS / RULES

P02

Intel 300mm InP Silicon Photonics: CPO Production Roadmap 2026

A03 | Bonding & assembly | Product adoption

Date not stated

FACT

Intel claims to be the only high-volume foundry capable of supplying 300mm indium phosphide (InP) silicon photonics at scale. This capability addresses scalability bottlenecks for next-generation near-package optics (NPO) and co-packaged optics (CPO) in AI and HPC applications.

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 21/25 | MATURITY 4/5 | DECISION LOCUS: TECH / PROCESS

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

P03

Photonics-Electronics Convergence: NEDO Project for AI Chiplets

A04 | Advanced package | Product adoption

Date not stated

FACT

NTT Innovative Devices leads a four-company consortium selected by NEDO to develop mass-production technologies for photonics-electronics convergence (PEC) devices. The project targets a tenfold increase in manufacturing productivity for optical chiplets to support high-capacity, low-power next-generation AI computing.

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 20/25 | MATURITY 4/5 | DECISION LOCUS: BUSINESS / RULES

P04

Samsung: TSMC 3nm process for Nvidia NVHBM base dies

A05 | Test & inspection | Research

Date not stated

FACT

Samsung Electronics will utilize TSMC's 3-nanometre process to produce base dies for Nvidia's next-generation custom high-bandwidth memory (HBM) chips. This marks the first time Samsung has collaborated with the Taiwanese foundry at the HBM manufacturing stage to meet Nvidia's standardization requirements.

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 1/5 | DECISION LOCUS: BUSINESS / RULES

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

P05

NVIDIA CPO Roadmap: OOI Prototype and 2028 Feynman Platform

A02 | Volume supply | Product adoption

Date not stated

FACT

NVIDIA has demonstrated a next-generation co-packaged optics (CPO) prototype featuring optics-on-interposer (OOI) and dense wavelength-division multiplexing (DWDM) to increase bandwidth density. While these specific technologies remain in the research stage, the company has already moved Spectrum-X Ethernet Photonics switches into full production for the Vera Rubin platform.

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 4/5 | DECISION LOCUS: SYSTEM / ADOPTION

P06

SK Hynix iHBM Cooling: HBM5 Thermal Resistance Reduction Specs

A10 | Materials | Product adoption

Date not stated

FACT

SK Hynix is collaborating with materials and equipment suppliers to commercialize iHBM cooling technology for HBM5 production. The initiative targets a reduction in thermal resistance of more than 30% to improve heat dissipation in next-generation memory.

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: SYSTEM / ADOPTION

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

P07

Unimicron IC Substrate Outlook: Equipment Lead Times and 2027 Gap

A07 | Bonding & assembly | Product adoption

Date not stated

FACT

Unimicron warns of a critical IC substrate crunch characterized by equipment lead times stretching to 30 months and a projected 20% material supply gap by 2027. These bottlenecks, driven by AI accelerator demand, threaten to desynchronize capacity expansion across the advanced packaging ecosystem.

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

Groq AI Chips: Shift to 3D-Stacked DRAM for Memory Capacity

A09 | Advanced package | Product adoption

Date not stated

FACT

Groq is considering a transition from SRAM-centric designs to 3D-stacked DRAM for its next-generation AI chips to address the widening gap between memory capacity and AI model growth. This shift aims to increase on-chip memory capacity while maintaining the high-speed performance required for inference workloads.

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

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

P09

Google TPU v9: Intel EMIB-T Packaging Yields and Timeline

A01 | Test & inspection | Research

Date not stated

FACT

Google is maintaining its commitment to Intel's EMIB-T packaging for the HumuFish TPU v9 variant despite reported substrate yields of 35% in 3Q26. While design revisions and tape-out delays may push shipments to 1Q28, the project serves as a critical test for Intel's advanced packaging capabilities.

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 1/5 | DECISION LOCUS: TECH / PROCESS

P10

Nikon DSP-110: Specs and Timeline for 1.5µm Digital Exposure

A12 | Volume supply | Research

Date not stated

FACT

Nikon will release the DSP-110 digital exposure system for advanced semiconductor packaging in April 2027, featuring a resolution of 1.5µm. The system increases throughput by approximately 40% compared to previous models and can be upgraded to 1.0µm resolution via optical system replacement.

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 1/5 | DECISION LOCUS: SYSTEM / ADOPTION

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 01-12

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	Google TPU v9: Intel EMIB-T Packaging Yields and Timeline	https://www.trendforce.com/news/2026/10/22/breaking-google-reportedly-sticks-with-intel-emib-t-despite-35-substrate-yields-as-humofish-schedule-tightens/
A02	NVIDIA CPO Roadmap: OOI Prototype and 2028 Feynman Platform	https://www.trendforce.com/news/2026/10/22/news-nvidia-unveils-next-gen-cpo-prototype-with-optics-on-silicon-interposer-targets-higher-bandwidth-density/
A03	Intel 300mm InP Silicon Photonics: CPO Production Roadmap 2026	https://www.trendforce.com/news/2026/10/20/news-intel-claims-industrys-only-high-volume-300mm-inp-silicon-photonics-capability-as-cpo-gains-momentum/
A04	Photonics-Electronics Convergence: NEDO Project for AI Chiplets	https://group.ntt/en/newsrelease/2026/10/20/261020a.html
A05	Samsung: TSMC 3nm process for Nvidia NVHBM base dies	https://www.thestar.com.my/tech/tech-news/2026/10/18/samsung-to-use-tsmc-technology-for-nvidia039s-next-generation-memory-chips-seoul-newspaper-reports
A06	NVIDIA Supply Chain: Capacity Expansion Roadmap for 2027-2028	https://nvidianews.nvidia.com/news/nvidia-and-supply-chain-partners-expand-capacity-to-accelerate-ai-infrastructure-buildout
A07	Unimicron IC Substrate Outlook: Equipment Lead Times and 2027 Gap	https://www.trendforce.com/news/2026/10/20/news-unimicron-warns-of-ic-substrate-crunch-up-to-30-month-equipment-lead-times-20-material-gap-in-2027/
A08	NVIDIA Vera CPU: Monolithic Die Architecture vs x86 Performance	https://www.trendforce.com/news/2026/10/19/news-nvidia-chooses-monolithic-die-over-chiplets-for-vera-cpu-claims-2x-x86-performance-in-agentic-ai/
A09	Groq AI Chips: Shift to 3D-Stacked DRAM for Memory Capacity	https://www.trendforce.com/news/2026/10/22/news-nvidia-linked-groq-reportedly-eyes-shift-from-sram-to-3d-stacked-dram-for-next-gen-ai-chips/
A10	SK Hynix iHBM Cooling: HBM5 Thermal Resistance Reduction Specs	https://www.digitimes.com/news/a20261022VL218/sk-hynix-cooling-production-materials-equipment.html
A11	Panel-Level Packaging: Scaling Challenges vs Wafer-Level Flow	https://semiengineering.com/what-breaks-when-packaging-moves-from-wafer-to-panel/
A12	Nikon DSP-110: Specs and Timeline for 1.5μm Digital Exposure	https://monoist.itmedia.co.jp/mn/articles/2610/19/news040.html

Semiconductor_BackEnd — Selected Articles

Date: 2026-10-25


Articles: 12

Table of Contents

- #01 Google TPU v9: Intel EMIB-T Packaging Yields and Timeline
- #02 NVIDIA CPO Roadmap: OOI Prototype and 2028 Feynman Platform
- #03 Intel 300mm InP Silicon Photonics: CPO Production Roadmap 2026
- #04 Photonics-Electronics Convergence: NEDO Project for AI Chiplets
- #05 Samsung: TSMC 3nm process for Nvidia NVHBM base dies
- #06 NVIDIA Supply Chain: Capacity Expansion Roadmap for 2027-2028
- #07 Unimicron IC Substrate Outlook: Equipment Lead Times and 2027 Gap
- #08 NVIDIA Vera CPU: Monolithic Die Architecture vs x86 Performance
- #09 Groq AI Chips: Shift to 3D-Stacked DRAM for Memory Capacity
- #10 SK Hynix iHBM Cooling: HBM5 Thermal Resistance Reduction Specs
- #11 Panel-Level Packaging: Scaling Challenges vs Wafer-Level Flow
- #12 Nikon DSP-110: Specs and Timeline for 1.5 μ m Digital Exposure

#01 Google TPU v9: Intel EMIB-T Packaging Yields and Timeline

Published October 22, 2026 TrendForce Taiwan



OVERVIEW

Google is maintaining its commitment to Intel's EMIB-T packaging for the HumuFish TPU v9 variant despite reported substrate yields of 35% in 3Q26. While design revisions and tape-out delays may push shipments to 1Q28, the project serves as a critical test for Intel's advanced packaging capabilities.

Key Findings

Google continues to use Intel's EMIB-T packaging for its HumuFish TPU v9 variant, despite reports that substrate yields were approximately 35% in 3Q26, missing the 50% year-end target. While there was speculation regarding a switch to TSMC's CoWoS-L, supply chain checks indicate HumuFish remains "100%" committed to EMIB-T. However, the project faces potential delays; the tape-out of the TSMC N2-based Compute Die A was pushed to late October, and design revisions could move the HumuFish tape-out to January 2027, potentially delaying shipments until 1Q28.

Technical Details

The yield challenges for EMIB-T substrates stem from two primary technical hurdles:

- Thermal expansion differences between ABF substrates and embedded silicon bridges, which complicate precise alignment during high-temperature processing.
- Air bubbles trapped in the Non-Conductive Film (NCF) applied beneath the bridges' TSVs during bonding, which can weaken adhesion.

To address these, Intel has adjusted its substrate electroplating process, while suppliers Ibiden, Unimicron, and Shinko are working to improve yields.

Background and Context

The TPU v9 lineup consists of three variants with split packaging orders:

- **WhaleFish:** Designed by Broadcom, using TSMC CoWoS-L, with mass production slated for 2H27.
- **HumuFish:** Designed by MediaTek, using Intel EMIB-T, with mass production targeted for 1H28.
- **TriggerFish:** An upgraded TPU v9 for stronger inference, using Intel EMIB-T, targeted for 2H28 or later. It features HBM4E (upgraded from HumuFish's HBM4), 2–3 times the SRAM capacity, and an additional simulation die.

Significance and Outlook

The success of the HumuFish timeline depends on how quickly substrate yields climb and whether design milestones are met. Because mass production is not slated until 1H28, the supply chain has a window to resolve EMIB-T issues, though a January 2027 tape-out would compress this cushion. The outcome represents a key test of Intel's ability to secure and maintain large external customers for its advanced packaging services.

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

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/>

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

Source: <https://www.trendforce.com/news/2026/10/22/breaking-google-reportedly-sticks-with-intel-emib-t-despite-35-substrate-yields-as-humufish-schedule-tightens/>

#02 NVIDIA CPO Roadmap: OOI Prototype and 2028 Feynman Platform

Published October 22, 2026 TrendForce Taiwan



OVERVIEW

NVIDIA has demonstrated a next-generation co-packaged optics (CPO) prototype featuring optics-on-interposer (OOI) and dense wavelength-division multiplexing (DWDM) to increase bandwidth density. While these specific technologies remain in the research stage, the company has already moved Spectrum-X Ethernet Photonics switches into full production for the Vera Rubin platform.

Key Findings

NVIDIA is developing next-generation CPO technology to bring optical engines closer to computing chips. At SEMICON West 2026, the company presented an optics-on-interposer (OOI) prototype and dense wavelength-division multiplexing (DWDM) technology. These research-stage developments are being evaluated for potential future products, though a commercialization timeline has not been announced.

Technical Details

The OOI design relocates optical engines from the organic package substrate directly onto the silicon interposer. By eliminating the organic substrate from the electrical signal path, signals travel through shorter, fine-pitch interposer connections, which is intended to reduce power consumption and increase bandwidth density. Additionally, NVIDIA demonstrated DWDM technology that utilizes eight wavelengths operating at up to 32 Gb/s each, achieving an aggregate bandwidth of 256 Gb/s per single optical fiber.

Background and Context

NVIDIA has already commercialized CPO through its Spectrum-X Ethernet Photonics switches, which entered full production in May 2026. These switches integrate 200Gb/s SerDes with CPO for the Vera Rubin platform and have been adopted by Oracle Cloud Infrastructure, Lambda, and CoreWeave. This production is supported by a broad ecosystem: TSMC handles silicon photonics fabrication, SPIL manages chip-scale packaging and testing, TFC packages laser dies, and Foxconn integrates the switches into networking systems. Other partners include Coherent, Lumentum, Sumitomo Electric, Corning, Senko, Browave, and Fabrinet.

Significance and Outlook

Future AI system scaling will depend on balancing circuit design with thermal management and packaging challenges. NVIDIA plans to implement optical interconnects in its Kyber NVL1152 architecture, which connects eight racks into a single 1,152-GPU NVLink domain. This architecture will support the Feynman platform, which is expected to arrive in 2028. TrendForce notes that the CPO/NPO market is expected to exceed US\$39 billion by 2030 as optical interconnects become critical for AI factory expansion.

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

Reference technology guide: Co-Packaged Optics (CPO): <https://troy-technical.com/semiconductor-back-end-term-cpo/>

Reference technology guide: Silicon Photonics: <https://troy-technical.com/semiconductor-back-end-term-siphotonics/>

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/>


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

Source: <https://www.trendforce.com/news/2026/10/22/news-nvidia-unveils-next-gen-cpo-prototype-with-optics-on-silicon-interposer-targets-higher-bandwidth-density/>

Collected: October 24, 2026 | Automated Research System (AI)

#03 Intel 300mm InP Silicon Photonics: CPO Production Roadmap 2026

Published October 20, 2026 TrendForce Taiwan



OVERVIEW

Intel claims to be the only high-volume foundry capable of supplying 300mm indium phosphide (InP) silicon photonics at scale. This capability addresses scalability bottlenecks for next-generation near-package optics (NPO) and co-packaged optics (CPO) in AI and HPC applications.

Key Findings

Intel has integrated InP-based laser technology into its 300mm silicon photonics manufacturing platform to enable high-volume production. This move targets the growing demand for high-bandwidth optical interconnects in AI and high-performance computing (HPC), specifically supporting NPO and CPO architectures.

Technical Details

Traditional InP substrates are typically limited to 150mm (6-inch) wafers, with 3-inch and 4-inch formats remaining common. InP crystals are characterized by brittleness, slow growth, and defect-control challenges, which historically limited laser chip output per wafer. By utilizing a 300mm platform—the standard for silicon-based CMOS manufacturing—Intel aims to improve production efficiency and scalability.

Industry Landscape and CPO Commercialization

Several major players are advancing CPO commercialization:

- **NVIDIA:** Shipping Spectrum-X CPO switches to select partners using TSMC's Compact Universal Photonic Engine (COUPE) technology.
- **Broadcom:** Shipping limited quantities of 51.2 Tb/s Bailly CPO switches and showcasing the third-generation Tomahawk 6-Davisson CPO switch, which offers 102.4 Tb/s capacity and 200 Gbps per channel.
- **TSMC:** Planning production of its COUPE-on-substrate solution in 2026, targeting a tenfold reduction in latency and twice the power efficiency of board-level pluggable optics.
- **Chinese Suppliers:** Zhongji Innolight has started small-volume production of 800G CPO optical modules, with shipments potentially beginning in 2027. Huawei introduced the NPO-based Ascend 960 supernode in September 2026.

Significance and Outlook

Despite manufacturing advances, testing remains a primary bottleneck for mass production. While optical alignment now takes seconds or less, testing a single photonic integrated circuit (PIC) can exceed 100 seconds. TrendForce identifies 2027 as the pivotal year for integrating CPO testing systems, with 2028 serving as the critical period for assessing large-scale production readiness.

Reference technology guide: Co-Packaged Optics (CPO): <https://troy-technical.com/semiconductor-back-end-term-cpo/>

Reference technology guide: Silicon Photonics: <https://troy-technical.com/semiconductor-back-end-term-siphotonics/>

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

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

Source: <https://www.trendforce.com/news/2026/10/20/news-intel-claims-industrys-only-high-volume-300mm-inp-silicon-photonics-capability-as-cpo-gains-momentum/>

#04 Photonics-Electronics Convergence: NEDO Project for AI Chiplets

Published October 20, 2026 NTT Group Japan



OVERVIEW

NTT Innovative Devices leads a four-company consortium selected by NEDO to develop mass-production technologies for photonics-electronics convergence (PEC) devices. The project targets a tenfold increase in manufacturing productivity for optical chiplets to support high-capacity, low-power next-generation AI computing.

Key Findings

A joint R&D proposal led by NTT Innovative Devices has been selected for NEDO's Research and Development Project of the Enhanced Infrastructures for Post-5G Information and Communication Systems. The consortium, comprising NTT Innovative Devices, Japan Semiconductor Corporation, Furukawa Electric Co., Ltd., and SHINKO ELECTRIC INDUSTRIES CO., LTD., will focus on establishing mass-production technologies for photonics-electronics convergence (PEC) devices. NTT will serve as an advisor for the project.

Technical Details

The project focuses on the development of optical chiplets, which are PEC devices that integrate photonic integrated circuits (PIC) and electronic integrated circuits (EIC) through high-density packaging. These chiplets replace electrical signals with optical signals for data transmission between processing chips, such as GPUs and CPUs, to overcome bandwidth limitations and power consumption increases associated with transmission distance.

The consortium aims to achieve the following technical milestones:

- Establishment of a pilot production line for future commercial manufacturing.
- A more than tenfold increase in manufacturing productivity compared to current R&D-stage processes.
- Development of the full manufacturing value chain, including photonic device fabrication, testing, optical interconnects, and advanced packaging.

Background & Context

The initiative responds to the unprecedented growth in data processing demand driven by generative AI in high-performance computing (HPC) systems and AI data centers. Since fiscal year 2023, NTT has led research into optical chiplet packaging technologies under a similar NEDO framework to address the rising power consumption of chip-to-chip communications.

Significance & Outlook

By combining expertise in silicon photonics, optical semiconductors, and semiconductor packaging, the participating companies intend to strengthen Japan's domestic semiconductor manufacturing capabilities. The project seeks to establish a globally competitive supply chain for PEC devices, ultimately reducing the energy consumption of AI-driven data centers and supporting a sustainable information and communications society.

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


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

Reference technology guide: Silicon Photonics: <https://troy-technical.com/semiconductor-back-end-term-siphotonics/>

Source: <https://group.ntt/en/newsrelease/2026/10/20/261020a.html>

#05 Samsung: TSMC 3nm process for Nvidia NVHBM base dies

Published October 18, 2026 The Star (Reuters) Malaysia



OVERVIEW

Samsung Electronics will utilize TSMC's 3-nanometre process to produce base dies for Nvidia's next-generation custom high-bandwidth memory (HBM) chips. This marks the first time Samsung has collaborated with the Taiwanese foundry at the HBM manufacturing stage to meet Nvidia's standardization requirements.

Key Findings

Samsung Electronics has agreed to use TSMC's 3-nanometre process for the production of base dies for Nvidia's next-generation custom high-bandwidth memory (HBM) chips. According to a report by the Seoul Economic Daily, this represents the first instance of Samsung working with the Taiwanese contract chipmaker during the HBM manufacturing stage.

Technical Details

Nvidia has specified a set of requirements for its NVHBM (Nvidia HBM) architecture, which was unveiled in August. The technical specifications requested from Samsung, SK Hynix, and Micron Technology include:

- **Base Dies:** Production via TSMC's 3-nanometre process.
- **Core Dies:** Implementation of sixth-generation DRAM technology.

The NVHBM design integrates Nvidia's custom memory controller directly into the base die. This architectural shift is intended to standardize implementation across multiple memory suppliers.

Background & Context

Samsung accepted these requirements following an internal review. However, the company intends to maintain its in-house production model for standard HBM products and custom orders for other clients. The move to TSMC is specific to the NVHBM requirements set by Nvidia.

Significance & Outlook

The collaboration highlights Nvidia's push for standardization across its memory supply chain by mandating specific foundry processes for the base die. While Samsung is currently adopting the TSMC process for these components, the company is also evaluating the feasibility of using its own foundry to produce NVHBM base dies should demand eventually exceed TSMC's available capacity.

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/>

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://www.thestar.com.my/tech/tech-news/2026/10/18/samsung-to-use-tsmc-technology-for-nvidia039s-next-generation-memory-chips-seoul-newspaper-reports>

Collected: October 24, 2026 | Automated Research System (AI)

#06 NVIDIA Supply Chain: Capacity Expansion Roadmap for 2027-2028

Published October 21, 2026 NVIDIA Newsroom USA



OVERVIEW

NVIDIA is partnering with key suppliers to aggressively expand manufacturing capacity for substrates, advanced packaging, testing, and optical engines to meet AI infrastructure demand. Key targets include a 2.5x increase in substrate capacity over three years and a 300% increase in advanced packaging capacity by 2028. These investments aim to resolve capacity constraints and strengthen global supply chain resiliency for AI supercomputing systems.

Key Findings

NVIDIA has established agreements with several supply chain partners to scale production of critical AI infrastructure components. The primary objective is to eliminate capacity constraints currently hindering the buildout of accelerated computing and full-stack AI factory platforms. Key expansions include a 2.5x increase in substrate capacity over the next three years and a projected doubling of total semiconductor testing capacity between calendar 2026 and 2027.

Technical Details

- **Substrates:** NVIDIA is providing a multibillion-dollar capacity reservation advance to IBIDEN for its Ono Plant in Japan, targeting a 3x increase in capacity with ramping starting in Q1 2028. Samsung Electro-Mechanics is adding capacity via existing operations in South Korea and new facilities in Sejong and Busan. Additionally, Unimicron's KF substrate expansion in Taiwan is expected to add capacity in Q2 2027, reaching full capacity by mid-2028.
- **Advanced Packaging:** NVIDIA and Amkor plan to increase interposer wafer capacity in South Korea by more than 300% by the end of 2028.
- **Testing:** A \$1.1 billion capacity reservation advance to KYEC will support new U.S. facilities, with high-volume production targeted for April 2027. Expanded operations in Singapore are expected to begin high-volume production in February 2027.
- **Optical Engines:** In collaboration with SPIL, NVIDIA is expanding manufacturing in Taiwan for co-packaged optics. Capacity is targeted to increase 15x by mid-2027, with a further doubling by the end of 2028.

Background & Context

According to NVIDIA founder and CEO Jensen Huang, the current AI infrastructure buildout is the largest in human history, creating extraordinary demand that has made capacity the primary constraint. To address this, NVIDIA is diversifying its supplier base and providing financial advances to ensure the scalability of production.

Significance & Outlook

These expansions across Japan, South Korea, Taiwan, Singapore, and the U.S. are designed to provide the redundancy and resiliency required to power the new industrial revolution. By securing these capacity increases through 2028, NVIDIA aims to ensure the availability of the hardware necessary for next-generation AI supercomputing platforms.

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

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

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

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

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

Reference company guide: Samsung Electro-Mechanics: <https://troy-technical.com/semiconductor-back-end-company-samsung-electro-mechanics/>

Source: <https://nvidianews.nvidia.com/news/nvidia-and-supply-chain-partners-expand-capacity-to-accelerate-ai-infrastructure-buildout>

#07 Unimicron IC Substrate Outlook: Equipment Lead Times and 2027 Gap

Published October 20, 2026 TrendForce Taiwan



OVERVIEW

Unimicron warns of a critical IC substrate crunch characterized by equipment lead times stretching to 30 months and a projected 20% material supply gap by 2027. These bottlenecks, driven by AI accelerator demand, threaten to desynchronize capacity expansion across the advanced packaging ecosystem.

Key Findings

Unimicron Chairman S.C. Chien reports significant supply chain constraints affecting IC substrates. Lead times for over 80% of key equipment have lengthened, with more than 20% of that equipment seeing delays exceeding double their original lead times; some lead times have reached 30 months. Regarding materials, a 20% supply shortfall is projected for 2027, specifically impacting ABF, copper-clad laminates (CCL), and low-CTE glass fiber cloth.

Technical Details

The technical demands on substrates are escalating as AI chip sizes grow. Substrate sizes expanded from roughly 1x reticle size in early designs to 3x in 2025, with targets of 5x for 2026 and 5–9x by 2027–2028, potentially reaching 12–14x. This scaling introduces several engineering hurdles:

- **Power and Signal Integrity:** AI chips can draw currents exceeding 1,000 amperes, increasing demands on thermal management, power integrity, and high-speed signal transmission.
- **Manufacturing Yields:** Stringent fine-line requirements for larger substrates may exceed the capabilities of individual makers. For example, Intel's EMIB substrate yields recently reached approximately 45%, with targets of 50% by 4Q26 and 60% by 1Q27.

Background and Context

These constraints emerge from a disparity in capacity expansion rates across the AI supply chain from 2024 to 2028. While hyperscaler capex is projected to grow 6.3-fold and CoWoS capacity 6.8-fold, IC substrate capacity is expected to merely double, and DRAM capacity is projected to expand 1.8-fold.

Significance and Outlook

As GPU-HBM platforms and custom ASICs evolve, IC substrates and PCBs are shifting from specification-based components to critical drivers of system-level performance. While different capacity bases mean growth multiples do not always equal shortages, the emergence of DRAM and substrate shortages in 2025 and 2026, combined with prolonged equipment lead times, complicates the industry's ability to synchronize capacity expansion.

[[END]]

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

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


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

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

Source: <https://www.trendforce.com/news/2026/10/20/news-unimicron-warns-of-ic-substrate-crunch-up-to-30-month-equipment-lead-times-20-material-gap-in-2027/>

#08 NVIDIA Vera CPU: Monolithic Die Architecture vs x86 Performance

Published October 19, 2026 TrendForce Taiwan



OVERVIEW

NVIDIA has opted for a monolithic compute die for its Vera CPU to optimize for agentic AI, claiming up to 2× the performance of leading x86 processors in sequential tasks. The architecture prioritizes low latency and high memory bandwidth to ensure GPUs remain highly utilized during complex AI agent tool calls.

Key Findings

NVIDIA's Vera CPU utilizes a monolithic compute die rather than the chiplet architectures employed by major x86 competitors. This design choice is specifically tailored for agentic AI workloads, which require strong single-threaded performance and low latency. According to reports, Vera delivers up to twice the performance of leading x86 processors for sequential tasks such as code compilation and Python scripting, with a 1.6-times advantage on a geometric-mean basis. Compared to its predecessor, Grace, Vera provides 1.5 times the performance.

Technical Details

The monolithic design eliminates the need for CPU cores to cross boundaries between separate chiplets, significantly reducing data movement latency. Key technical specifications include:

- **Bisection Bandwidth:** 3.4 TB/s, which is three times that of the leading x86 processor used in NVIDIA's comparison.
- **Memory Subsystem:** A custom LPDDR5X memory subsystem.
- **Interconnect:** Second-generation Scalable Coherent Fabric.
- **Spatial Multithreading:** A technology that allocates dedicated resources to I/O and worker threads to prevent background operations from interfering with primary processing tasks.

Background and Context

The shift in architecture reflects a change in CPU design priorities. While traditional cloud workloads focus on throughput and high core counts, agentic AI requires the ability to execute sequential tasks quickly. Consequently, NVIDIA is expanding the role of the CPU in AI infrastructure. While the Grace CPU was primarily a host processor for GPU systems, the demand for reinforcement learning and agentic AI has led NVIDIA to develop standalone Vera CPU racks to handle CPU-intensive tasks and AI agent tool calls.

Additional Updates

Separately, NVIDIA Senior Vice President Jeff Fisher has denied rumors regarding the discontinuation of the GeForce RTX 5090. Despite global supply constraints and market prices exceeding the US\$1,999 MSRP on some platforms, Fisher stated the RTX 5090 remains a key product that the company will continue to supply.

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


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

Source: <https://www.trendforce.com/news/2026/10/19/news-nvidia-chooses-monolithic-die-over-chiplets-for-vera-cpu-claims-2x-x86-performance-in-agentic-ai/>

Collected: October 24, 2026 | Automated Research System (AI)

#09 Groq AI Chips: Shift to 3D-Stacked DRAM for Memory Capacity

Published October 22, 2026 TrendForce Taiwan



OVERVIEW

Groq is considering a transition from SRAM-centric designs to 3D-stacked DRAM for its next-generation AI chips to address the widening gap between memory capacity and AI model growth. This shift aims to increase on-chip memory capacity while maintaining the high-speed performance required for inference workloads.

Key Findings

Groq, the AI chip company that licensed its inference technology to NVIDIA, is exploring the adoption of 3D-stacked DRAM (3DS) for its next-generation chips. This potential move away from a primary reliance on SRAM is driven by the need for greater memory capacity as AI workloads expand beyond the limits of traditional SRAM architectures.

Technical Details

The proposed 3DS memory utilizes through-silicon vias (TSVs) to connect DRAM dies vertically. While both 3DS and high-bandwidth memory (HBM) utilize TSVs, 3DS specifically prioritizes increasing memory capacity while retaining conventional DRAM interfaces, whereas HBM focuses on data transfer speeds.

Current specifications for NVIDIA's Groq 3 LPX, designed for the Vera Rubin platform, include:

- Per LPU: 500MB of SRAM with 150TB/s on-chip memory bandwidth.
- Per LPX rack (256 LPUs): 128GB of SRAM and 40PB/s aggregate SRAM bandwidth.
- Additional capacity: Each rack provides 12TB of DDR5 memory alongside the SRAM.

Background & Context

The shift is necessitated by a disparity in scaling: SRAM capacity increases by approximately 10–20% every two years, while AI model sizes can double every six to nine months. This makes it increasingly difficult for SRAM alone to meet demand. Other industry players are pursuing similar paths; Cerebras is developing the CS-6 system, which will integrate stacked DRAM with wafer-scale SRAM and compute.

Significance & Outlook

Manufacturing competition for these next-generation components is intensifying. Samsung Foundry currently produces the Groq 3 LPU using a 4nm process at its Pyeongtaek facility, and Samsung executives have discussed further foundry cooperation with NVIDIA. Simultaneously, TSMC has stated it is working with customers on the development of next-generation LPUs.

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

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

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

Source: <https://www.trendforce.com/news/2026/10/22/news-nvidia-linked-groq-reportedly-eyes-shift-from-sram-to-3d-stacked-dram-for-next-gen-ai-chips/>

Collected: October 24, 2026 | Automated Research System (AI)

#10 SK Hynix iHBM Cooling: HBM5 Thermal Resistance Reduction Specs

Published October 23, 2026 DIGITIMES Taiwan



OVERVIEW

SK Hynix is collaborating with materials and equipment suppliers to commercialize iHBM cooling technology for HBM5 production. The initiative targets a reduction in thermal resistance of more than 30% to improve heat dissipation in next-generation memory.

Key Findings

SK Hynix is reportedly engaging with its supply chain of materials and equipment providers to integrate a next-generation cooling technology known as iHBM into the production of HBM5. The primary objective of this collaboration is to achieve a significant improvement in thermal management for high-bandwidth memory.

Technical Details

The technical focus of the iHBM cooling technology is the reduction of thermal resistance. According to the report, SK Hynix is targeting a reduction of more than 30% in thermal resistance compared to previous standards. By working with specialized suppliers, the company intends to commercialize these improvements to ensure that HBM5 can operate within necessary thermal envelopes despite increasing performance demands.

Background & Context

As memory architectures evolve toward HBM5, managing heat becomes a critical bottleneck for performance and reliability. The transition to iHBM cooling represents a strategic shift toward optimizing the physical materials and equipment used during the manufacturing process to enhance heat dissipation efficiency.

Significance & Outlook

The successful commercialization of a 30% more efficient cooling solution would address one of the most pressing challenges in high-performance computing memory. By securing the necessary materials and equipment through supplier partnerships, SK Hynix is positioning its HBM5 production to handle higher thermal loads, which is essential for the stability of next-generation AI and computing hardware.

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


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

Source: <https://www.digitimes.com/news/a20261022VL218/sk-hynix-cooling-production-materials-equipment.html>

Collected: October 24, 2026 | Automated Research System (AI)

#11 Panel-Level Packaging: Scaling Challenges vs Wafer-Level Flow

Published October 22, 2026 Semiconductor Engineering USA



OVERVIEW

Transitioning advanced packaging from 300mm wafers to larger panels significantly increases unit yield for large AI and HPC packages, with some configurations producing 8 times more units than wafers. However, this shift requires a complete reconsideration of equipment and processes, as traditional wafer-based radial symmetry and handling systems do not transfer to rectangular panels.

Key Findings

Moving from wafer-level to panel-level packaging dramatically improves area utilization, particularly for large rectangular AI and HPC packages. For a 5.5×-reticle-size interposer, a 300mm wafer produces 9 units, whereas a 600mm panel produces 72 units—an eightfold increase in utilization. Intel Foundry demonstrated a similar advantage with a 4× reticle package, producing approximately 48 units on a 510 × 515mm panel compared to 10 units on a 300mm wafer.

Technical Details

The transition introduces significant mechanical and process challenges:

- **Geometry and Coating:** Traditional spin coating relies on radial symmetry, which causes uneven film thickness on rectangular panels. Consequently, 600mm panel flows utilize slit or slot-die coating across active areas of roughly 586 × 586mm.
- **Lithography:** Conventional wafer steppers are being displaced by maskless direct imaging in many panel flows.
- **Process Control:** While wafer processing is typically single-sided, panel processing is dual-sided, complicating uniformity and Cpk (process capability index).
- **Precision:** As line and space dimensions move from 5µm toward 2µm and below, panel processes require precision closer to front-end manufacturing.

Background & Context

The industry has spent decades optimizing equipment and materials for round 300mm wafers. Because the largest square that fits on a 300mm wafer is only about 210 × 210mm, panels have become an obvious choice for the scaling needs of AI systems. However, the manufacturing infrastructure is not yet fully ready for the larger format.

Significance & Outlook

Due to high risks in assembly on 600mm panels, some manufacturers are adopting a compromise: completing the interposer in a large 600mm format and then dividing it into four 300mm pieces for assembly and the C4 process. This approach yields 64 units—fewer than a full intact panel but significantly more than a wafer. This transition reveals that many "process requirements" were actually dependencies on wafer geometry, necessitating a new manufacturing process rather than a simple scale-up.

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

Reference technology guide: Redistribution Layer (RDL): <https://troy-technical.com/semiconductor-back-end-term-rdl/>

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

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


Reference company guide: ASE Group: <https://troy-technical.com/semiconductor-back-end-company-ase-group/>

Source: <https://semiengineering.com/what-breaks-when-packaging-moves-from-wafer-to-panel/>

Collected: October 24, 2026 | Automated Research System (AI)

#12 Nikon DSP-110: Specs and Timeline for 1.5 μ m Digital Exposure

Published October 20, 2026 MONOist (ITmedia) Japan



OVERVIEW

Nikon will release the DSP-110 digital exposure system for advanced semiconductor packaging in April 2027, featuring a resolution of 1.5 μ m. The system increases throughput by approximately 40% compared to previous models and can be upgraded to 1.0 μ m resolution via optical system replacement.

Key Findings

Nikon announced the upcoming launch of the DSP-110, a digital exposure system designed for semiconductor advanced packaging. Scheduled for release in April 2027, the device provides a resolution of 1.5 μ m (L/S). A primary technical advantage is its flexibility; the system can be modified to achieve a 1.0 μ m resolution by replacing the optical system to meet future process requirements.

Technical Details

The DSP-110 utilizes an i-line equivalent light source and achieves a throughput of 72 panels per hour for 510 $\times$ 515mm substrates. By redesigning the optical system for 1.5 μ m resolution, Nikon has improved throughput by approximately 40% compared to the DSP-100. The system maintains an overlay accuracy of $\pm 0.3\mu$ m or less.

The device supports a wide range of substrate sizes, including:

- Square substrates: 310 $\times$ 310mm, 510 $\times$ 515mm, 600 $\times$ 600mm, and 730 $\times$ 920mm
- Wafers: 300mm

Additionally, the redesigned optical system expands the focus margin, which is intended to stabilize patterning across diverse substrate conditions and improve yield.

Background and Context

The development of the DSP-110 responds to the growing demand for high-density wiring in multi-layer FC-BGA substrates and large interposers within the advanced packaging market. This has increased the importance of Redistribution Layer (RDL) exposure technology. Nikon's strategy involves providing digital exposure solutions that eliminate the need for photomasks while balancing high resolution with productivity. The DSP-110 follows the previously released DSP-100, which offers 1.0 μ m resolution and supports 600mm square substrates.

Significance and Outlook

The DSP-110 is positioned to reduce total costs and increase production efficiency for manufacturers. Because the system allows for a resolution upgrade to 1.0μm, it provides a roadmap for users to adapt to the increasing resolution requirements expected as advanced packaging technology evolves.

Reference technology guide: Redistribution Layer (RDL): <https://troy-technical.com/semiconductor-back-end-term-rdl/>

Source: <https://monoist.itmedia.co.jp/mn/articles/2610/19/news040.html>

Collected: October 24, 2026 | Automated Research System (AI)