intel foundry

Platform Brief

Accelerating AI and HPC with advanced process and packaging technologies

Intel Foundry offers a unique portfolio of process nodes, technologies, and packaging solutions to help AI and HPC silicon companies scale their performance.

Contents

AI and HPC demand high-performance silicon

The explosive growth of generative and agentic AI is driving record demand for powerful, energy-efficient computing solutions. This has transformed semiconductor manufacturing to chiplet-based architectures, enabling the modular assembly of dies based on different nodes from various vendors, thereby boosting performance and reducing costs. Simultaneously, supply chains are being restructured to enhance resilience, mitigate geopolitical risk, and promote sustainability.

Meeting these demands requires foundries to deliver advanced technologies, including leading-edge nodes for optimal performance per watt, robust multi-die interconnects, high-bandwidth memory solutions, and efficient power delivery and thermal management for reliability. Supporting large, customizable multi-die systems provides designers with the flexibility and integration required for next-generation AI workloads.

Intel Foundry delivers a comprehensive portfolio of technologies that empower advanced AI and high-performance computing (HPC) designs, combining cutting-edge process nodes, high-performance packaging, and a broad intellectual property (IP) ecosystem.

Intel's portfolio of process nodes for AI and HPC

Intel Foundry delivers a portfolio of advanced process nodes designed to meet the demanding performance, power, and area (PPA)targets of AI and HPC workloads. Starting with Intel 18A, these nodes integrate breakthrough technologies for compute density, power delivery, and performance.

Intel 18A

Intel 18A introduces **RibbonFET** gate-all-around (GAA) transistors and **PowerVia** backside power delivery, offering significant performance improvements (see Figure 1). These technologies enable higher transistor performance, improved electrostatic control, and reduced IR drop, critical for sustaining AI and HPC workloads under extreme power and thermal constraints. Researched, designed, and produced in the U.S., Intel 18A will be used to manufacture next-generation Intel® Core™ Ultra and Xeon® processors.

Intel 18A-P extends these capabilities with optimized transistor ribbon sizes and transistor threshold voltage options.

Both nodes support advanced standard cell libraries, **high-performance (180H)** and **high-density (160H)**, which achieve over **30% density improvement** compared to Intel 3. PowerVia frees up front-side routing for signals, enabling tighter layouts that reduce resistance-capacitance (RC) delay, while Electronic Design Automation (EDA) tool support ensures seamless integration.

Intel 18A delivers		Intel 18A-P delivers
>15% performance per watt increase compared to Intel 3	**1.3x** chip density compared to Intel 3	**8%** performance per watt increase compared to Intel 18A

Figure 1. Intel 18A delivers improvements in power efficiency and chip density compared to Intel 3.[1]

RibbonFET

RibbonFET is Intel Foundry's first new transistor architecture since FinFET in 2011, introducing GAA with ribbon-shaped channels fully surrounded by the gate. This structure provides superior electrostatic control, enabling stable operation at lower voltages and reducing leakage—critical for power-constrained AI accelerators. A single ribbon stack delivers the same current as multiple FinFET fins in a smaller footprint, boosting density and supporting massive parallelism in compute arrays. RibbonFET also improves short-channel control for high-speed logic and scales efficiently for future nodes.

PowerVia

PowerVia is the industry's first backside power delivery technology, now supported by EDA reference flows for Intel 18A designs. Relocating the power grid to the die's backside improves signal integrity and frees front-side metal for logic, enabling **up to 10% higher cell utilization** and **up to 4% performance gain at constant power.**[2] For AI and HPC workloads, this architecture is critical, as it reduces IR drop and voltage droop, ensuring stable operational voltage and minimizing timing violations during rapid, high-current transients that can reach hundreds of amps in large compute arrays. Cost impact is offset by streamlined lower-metal patterning, making PowerVia a practical and future-ready solution for advanced AI and HPC architectures.

Intel Foundry partnered with EDA ecosystem vendors to optimize design flows for PowerVia, including improved routing engines, clock network shielding, buffer placement, and a reduction in buffer count.

Omni MIM

Omni MIM provides the industry's highest announced capacitance density (397 fF/μm²) for metal-insulator-metal capacitors, embedded in backside metals and adaptable for advanced packaging. This allows high-density decoupling on the chiplet, reducing inductance and improving transient response for AI and HPC systems. See 'Spotlight on power delivery' for more on our power delivery solution.

Intel 18A-PT for 3D Integration

Intel 18A-PT is the industry's **first announced base die with backside power delivery**, designed for next-generation 3D integrated circuit (3DIC) architectures (see Figure 2). Building on 18A-P, it adds **pass-through and die-to-die through-silicon vias (TSVs)** along with hybrid bonding, enabling both logic-on-logic and memory-on-logic stacking to bring high-bandwidth memory closer to compute resources. This proximity reduces latency and increases memory density, critical for AI workloads that require rapid data access and large model storage. Risk production is planned for 2028.

Intel 18A-PT delivers		
~9x die-to-die bandwidth density compared to Intel 3-T	**20-25%** chip density increase compared to Intel 3-T	**25-35%** power reduction compared to Intel 3-T

Figure 2. Intel 18A-PT improves bandwidth density, chip density, and power efficiency compared to Intel 3-T (Intel Foundry's base die derivative of Intel 3), offering a more capable base die for multi-chip packages.[3]

Figure 3. RibbonFET (left) introduces gate-all-around transistors for improved control, higher drive current, and a smaller footprint. PowerVia (right) delivers backside power delivery to enhance chip scaling, reduce IR drop, and simplify lower metal patterning.

Intel 14A

Building on Intel 18A, Intel 14A introduces **RibbonFET 2, the second generation of gate-all-around (GAA) transistors,** and **PowerDirect, an advanced backside power delivery technology**, delivering higher performance per watt and greater transistor density (see Figure 4). Intel 14A-E adds targeted enhancements to improve power efficiency.

A major innovation is **Turbo Cells**, specialized libraries for performance-critical paths. Double-height Turbo Cells provide exceptional drive current in an area-efficient format, enabling higher CPU frequencies and GPU throughput with minimal power impact. Integrated into EDA flows, Turbo Cells enable designers to balance high-performance and power-efficient cells within compute-intensive blocks.

Intel 14A is set to deliver		
15-20% performance per watt increase (at the same power) compared to Intel 18A	**1.3x** chip density compared to Intel 18A	**25-35%** power reduction (at the same performance) compared to Intel 18A

Figure 4. Intel 14A improves performance, chip density, and power efficiency compared to Intel 18A.[4]

Design, packaging, and test services for AI and HPC

As die sizes increase, AI and HPC developers require advanced packaging technologies to integrate high-density compute with memory that achieves high yields. Intel Foundry Advanced System Assembly & Test (Intel Foundry ASAT) leads in chiplet technology, offering end-to-end packaging and test solutions that accelerate cycle times for complex multi-chip systems and improve yield.

With over a decade of chiplet design and production experience, we enable flexible integration, mixing chiplets from any foundry to combine Intel dies with third-party components or delivering packaging-only services. For large, complex solutions, Embedded Multi-Die Interconnect Bridge (EMIB) technology enhances die-to-die connectivity, enabling systems that exceed 8 times the reticle size.

Our portfolio spans architectural optimization, layout, and analysis for cost, performance, and yield, plus test-only services for sorting and system-level validation. We're advancing next-generation technologies, such as glass substrates and co-packaged optics, which promise transformative bandwidth and scalability for future AI and HPC systems.

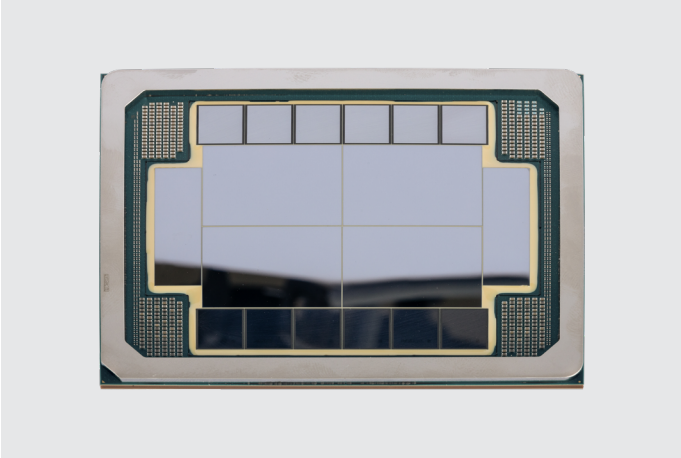

Figure 5. An AI chip test vehicle showing an architecture with high-bandwidth memory at the image top and bottom, high-speed communications at the left and right of the image, and GPU space in the middle. A test vehicle tests the manufacturing process but does not create a working product.

Embedded Multi-Die Interconnect Bridge – EMIB 2.5D

Intel Foundry's unique EMIB 2.5D interconnect supports high-data-rate signaling between adjacent dies, utilizing simple driver/receiver circuitry to meet the high data throughput requirements of AI and HPC applications. It embeds die-to-die bridges in the substrate, eliminating the need for and cost of an interposer and the associated wafer-level assembly stacking process. This cuts manufacturing time and eliminates any wafer assembly yield loss with the ability to increase the area of the package.

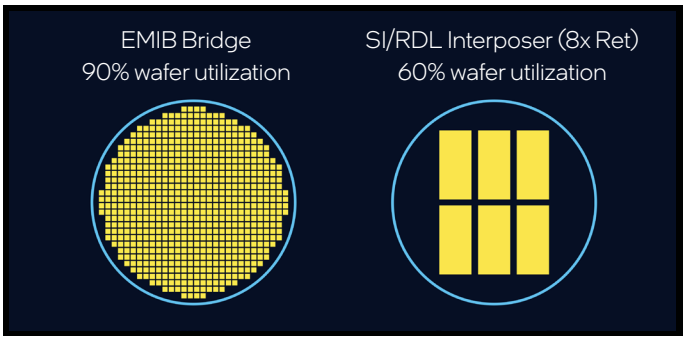

Figure 6. EMIB offers higher wafer utilization and yield compared to interposer architectures.

Our latest version, EMIB-T 2.5D, adds TSVs for stitching the die to HBM4 memory and ultra-high-speed UCIe die-to-die communications. EMIB-T enables designs to be converted from other packaging technologies with minimal redesign. EMIB-T also enables vertical power delivery to minimize DC and AC noise.

EMIB 2.5D supports the growth of AI and HPC solutions with more than 8 times reticle size in 2026 and 12 times in 2028 (see Figure 7).

Year	Reticle size	Package size	HBMs	EMIBs
2023	~4x	~80x80	8	12
2026	>8x	~120x120	12	>20
2028	>12x	>120x180	>24	>38

Figure 7. EMIB delivers scalability now and in the future to support AI demand.

EMIB 3.5D and Foveros Direct 3D

EMIB 3.5D combines embedded bridges with vertical die-stacking to deliver high performance and flexibility. Foveros 2.5D stacks chiplets on a passive base die, while Foveros Direct 3D stacks them on an active base die, increasing functionality. Foveros Direct 3D uses Hybrid Bonding Interconnect, with sub-5µm copper-to-copper bump pitches, to ensure ultra-low power, low-resistance connections.

Advanced Chiplet Test

As package sizes grow and complexity increases, comprehensive end-to-end testing is essential for quality and cost control. With final products often worth thousands, identifying known good dies is critical for maximizing yield. Intel Foundry's singulated die sort test accelerates the identification of known good dies. Our advanced methodology uses active thermal controls for precise, consistent testing.

Additionally, our system-level testing (SLT) detects marginal chip defects missed by traditional methods. SLT simulates real-world scenarios to uncover issues that only appear under actual operating conditions.

Our end-to-end testing solutions are optimized for our proprietary High Density Modular Test (HDMT) equipment. We also offer testing on commercial automatic testing equipment (ATE) to meet your needs. With these solutions, Intel Foundry delivers reliable, high-quality products.

Spotlight on power delivery

Generative and agentic AI workloads exert intense power demands on multi-die systems, causing rapid current surges at microsecond and nanosecond timescales that can destabilize conventional power delivery networks.

Intel Foundry solves these challenges with a holistic stack of innovations, beginning at the transistor level. Intel 18A logic tiles feature PowerVia backside power delivery, which dramatically shortens the power path, reduces IR drop, and frees up frontside routing for high-speed signaling. Paired with on-die Omni MIM capacitors, this architecture minimizes inductive droop during fast transients, helping to maintain frequency during bursty AI loads. Omni MIM capacitors also shrink power delivery network (PDN) impedance at high frequencies.

At the package level, EMIB-T 2.5D incorporates through-silicon vias (TSVs) and high-power MIM capacitors directly in the bridge, eliminating cantilevered power traces and enabling vertical power delivery with lower impedance and improved reliability. These improvements are essential for meeting the increased bandwidth requirements of HBM4 and future generations, as well as UCIe operating at speeds of 32Gbps and higher.

To further stabilize power delivery, Intel Foundry integrates embedded deep trench capacitors (eDTC) for ultra-high capacitance density close to the load, thereby flattening PDN impedance across the base die. Embedded MIM (eMIM-T) technology delivers superior decoupling at the base die, leveraging the same advanced architecture as Omni MIM even on mature nodes. For precise, high-speed point-of-load regulation, CoaxMIL technology embeds magnetic inductors into the substrate, enabling integrated voltage regulators (IVRs) with rapid response and minimal loss.

Together, these technologies form a scalable power delivery fabric that keeps pace with rising compute density.

eMIM-T
>2.5 µF/mm^2

CoaxMIL
2.5-5.0 nF/mm^2
>10 A/mm^2

eDTC
>2.5 µF/mm^2

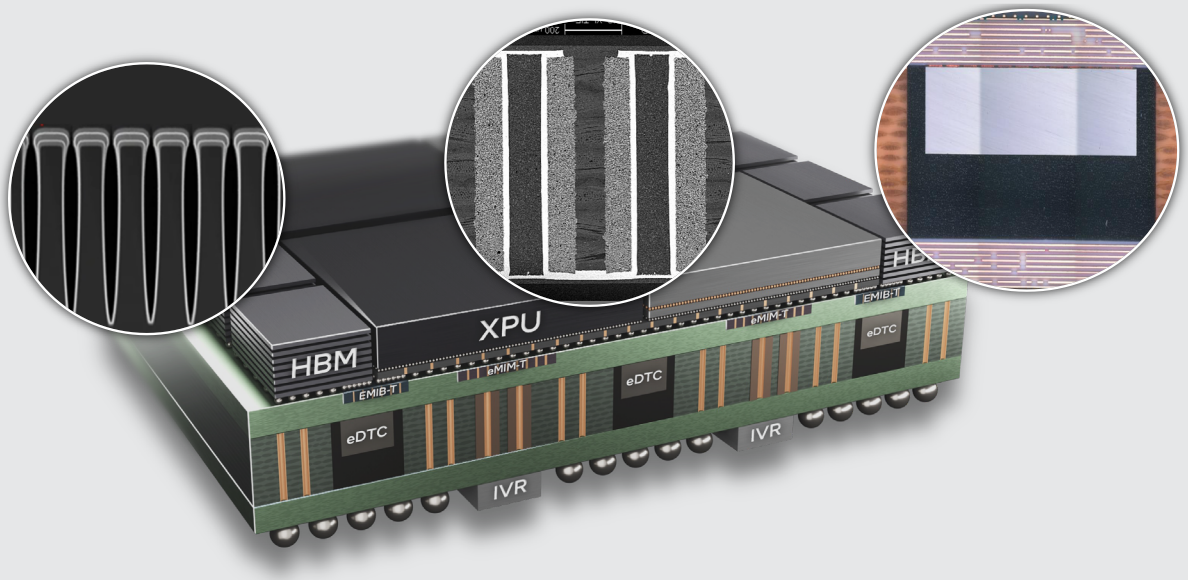

Figure 8. An example of Intel Foundry's power delivery solution shown at Intel Direct Connect 2025, showcasing forward-looking plans for increased eMIM-T capacity. All product plans and roadmaps subject to change.

Tap into the Intel Foundry ecosystem

The Intel Foundry Accelerator Alliance helps customers turn their innovations into silicon by leveraging the world-class capabilities of partners from across the design ecosystem, including IP, proven tools and flows, cloud capabilities, and design acceleration services. Recent alliance expansions add collaboration for chiplet interoperability and streamlined spec-to-silicon delivery, potentially accelerating complex application time to market. These efforts complement co-optimization initiatives that align process and system technologies with workload-specific needs.

To meet the demands for high-bandwidth, low-latency die-to-die communication from AI and HPC, Intel Foundry supports **UCIe 1.1 and 2.0**, the open standards for chiplet interconnect. UCIe ensures interoperability across foundries, reduces integration cost, and speeds deployment. UCIe 2.0 introduces optical interconnect and advanced packaging support, enabling modular and scalable architectures for next-generation AI and HPC systems.

Transform your silicon vision into reality with Intel Foundry

Intel Foundry offers full-stack solutions that accelerate time to market, energizing the industry transition from "system on chip" to "system of chips." Our integrated approach combines cutting-edge process node technologies—from mature nodes to leading-edge processes—with revolutionary advanced packaging solutions and select systems technologies.

Ready to revolutionize your next-generation products? Visit intel.com/foundry.

intel foundry

[1] Based on Intel internal analysis comparing Intel 18A to Intel 3 as of February 2024. Results may vary.

[2] Hafez, et. al. "Intel PowerVia Technology: Backside Power Delivery for High Density and High-Performance Computing." June 2023.

[3] Based on Intel internal analysis comparing Intel 18A-PT to Intel 3-T as of April 2025. Results may vary.

[4] Based on Intel internal analysis comparing Intel 14A to Intel 18A as of April 2025. Results may vary.

1125/JM/CAT/PDF 367370-001US