

# Generative AI Hardware Materials Market 2026-2036-Semiconductors, Memory, Packaging, and Thermal Management

<https://marketpublishers.com/r/GA36C20C2884EN.html>

Date: May 2026

Pages: 438

Price: US\$ 1,600.00 (Single User License)

ID: GA36C20C2884EN

## Abstracts

Generative AI has become the largest single demand driver in the semiconductor industry, and the Generative AI Hardware Materials market is the supply-side response to that demand. It spans the silicon, memory, packaging, photonics, thermal, and power-delivery layers that go into AI infrastructure across hyperscale data centres, enterprise and neocloud deployments, sovereign-AI programs, and the emerging edge AI tier.

The market is best understood as nine concentric layers of the AI compute stack. AI accelerator silicon sits at the top — GPUs from NVIDIA and AMD, custom hyperscaler ASICs from Google, AWS, Microsoft, and Meta, and challenger architectures from Cerebras, Groq, SambaNova, and the Chinese sovereign-AI silicon cohort. Beneath the accelerator die sits high-bandwidth memory, which has emerged as the most valuable layer below the compute silicon and the principal beneficiary of the HBM3E-to-HBM4-to-HBM5 roadmap. Advanced 2.5D and 3D packaging — CoWoS, SoIC, and the emerging glass-core substrate ecosystem — integrates compute and memory dies into the physical packages that AI accelerators ship in. Co-packaged optics and silicon photonics are moving from pilot to volume as electrical signalling reaches its limit above 224 Gbps per lane. Thermal management is shifting from air cooling to direct-to-chip liquid cooling, immersion, and in-package microfluidic cooling as accelerator TDPs scale past 1500 W. Power delivery is transitioning from 12V to 48V to 800V HVDC architectures, pulling GaN and SiC into data centre PSU applications. Networking silicon and optical components, the data centre construction supply chain, and the edge AI silicon tier round out the stack.

Frontier-model performance is now bounded by physical limits that yield only to materials and packaging innovation — compute throughput by reticle area and transistor

density, memory bandwidth by HBM stack height and pin width, interconnect bandwidth by copper trace attenuation, thermal dissipation by TIM conductivity and coolant flow rate, and power delivery by IR drop and voltage-regulator efficiency. Each of these walls is being attacked by a specific materials or packaging innovation, creating a sustained, multi-layer demand expansion across the supply chain.

The supply base is structurally Asia-centric. Taiwan dominates leading-edge logic and advanced packaging, Korea dominates HBM, Japan dominates specialty materials and substrate inputs, and China is building a parallel sovereign-AI hardware stack under export-control constraints. The materials and packaging layer of the GenAI supply chain is one of the most concentrated industrial value chains in the modern economy, and its trajectory will define the cadence at which AI compute scales over the next decade.

The Generative AI Hardware Materials Market 2026–2036 is the most comprehensive single source on the materials- and packaging-layer supply side of the generative AI hardware build-out. It complements demand-side coverage of foundation models, AI services, and hyperscaler capex by quantifying the physical infrastructure — silicon dies, HBM stacks, advanced packages, substrates, photonics, thermal systems, and power semiconductors — that hyperscaler AI capex commitments translate into across the supply chain.

The report covers nine concentric layers of the AI hardware materials value chain in dedicated chapters: AI accelerator silicon, AI-driven chip design (EDA), high-bandwidth memory and beyond-HBM architectures, advanced packaging and substrates, co-packaged optics and silicon photonics, thermal management, power delivery and the GaN/SiC transition, networking and optical materials, the data centre construction supply chain, and the edge GenAI hardware tier. Each chapter combines bottom-up unit-volume and ASP analysis, capacity and capex tracking, technology-roadmap mapping, and detailed company profiles. Regional analysis covers Taiwan, South Korea, Japan, China, Southeast Asia and India, the United States, Europe, and Israel. A dedicated supply-chain and geopolitics chapter covers the US-China technology competition, Taiwan concentration risk, critical-materials supply, CHIPS Act and European Chips Act implementation, and the parallel China sovereign-AI hardware stack. Sustainability and embodied-carbon analysis covers the operational and embodied emissions profile of AI infrastructure, the PFAS chemistry transition, and the carbon-accounting regulatory framework.

The methodology aggregates segment-level forecasts built from bottom-up unit volumes, ASPs, and content-per-unit analysis, with Base, Bull, and Bear scenarios

through 2036 and regional capture forecasts for nine geographies. The strategic outlook frames five defining themes of the GenAI hardware decade, a choke-point map of binding constraints, a strategic investment framework, and an M&A landscape analysis through 2030.

The report is designed for buyers and decision-makers in the Asian foundry, OSAT, memory, substrate, photonics, thermal, and cooling vendor ecosystem; for hyperscalers and AI silicon designers evaluating capacity and supplier strategy; for institutional investors building positions across the AI hardware value chain; and for sovereign-AI program managers planning national AI infrastructure. Coverage spans the full decade from 2026 through 2036 with dedicated treatment of the major architectural inflections, capacity bottlenecks, technology transitions, and geopolitical scenarios that will define the GenAI hardware decade. The result is a single integrated source on the hardware that makes generative AI physically possible.

### **Contents include:**

Executive Summary — Key findings; the GenAI hardware bottleneck; materials value chain at a glance; ten-year forecast highlights; strategic implications for Asian foundries, OSAT, memory, substrate, and cooling vendors; major market players (NVIDIA, TSMC, SK hynix, Samsung Electronics, ASE Technology)

The Compute Stack Behind Generative AI — Training vs. inference economics; pre-training, post-training, RLHF compute splits; inference token economics and serving infrastructure; test-time compute and reasoning-model demand; cloud, edge, and sovereign AI; hyperscaler clusters at 100,000-GPU scale; enterprise on-prem and neocloud deployments; sovereign AI build-outs; the memory wall in LLM serving; HBM ASP as percentage of AI accelerator BOM; CoWoS as the constraining bottleneck; hyperscaler vs. enterprise vs. sovereign capex

AI Accelerator Silicon — NVIDIA Hopper ? Blackwell ? Blackwell Ultra ? Rubin ? Rubin Ultra roadmap; NVL72 rack architecture and post-Rubin scale-up; AMD MI300X ? MI355X ? MI400 trajectory; Intel Gaudi and post-Gaudi; custom hyperscaler ASICs (Google TPU, AWS Trainium/Inferentia, Microsoft Maia/Cobalt, Meta MTIA); ASIC NRE economics and break-even analysis; domain-specific architectures (Cerebras WSE-3, Groq LPU, SambaNova RDU); Chinese AI chip ecosystem (Huawei Ascend, Cambricon, Biren, Moore Threads); TSMC, Samsung, Intel, and SMIC process-node roadmaps; EUV and High-NA EUV adoption; wafer-level integration and reticle stitching

**AI-Driven Chip Design (EDA)** — The EDA bottleneck in the AI hardware era; the recursive loop of AI designing AI hardware; incumbent EDA vendors' AI initiatives; the startup cohort across agentic AI for digital design and verification, physics-AI for simulation and advanced packaging, AI for analog and PCB design, and EDA-adjacent silicon; geographic distribution; AI-EDA tools market forecast 2026–2036

**High Bandwidth Memory and Beyond** — HBM architecture and TSV stacking fundamentals; HBM3/HBM3E, HBM4/HBM4E, HBM5/HBM5E generation roadmap; SK hynix, Samsung, and Micron strategy and capacity outlook; bit-shipment and wafer-capacity forecasts; custom HBM (cHBM) and base-die innovation; standard vs custom HBM revenue split; compute-in-memory and processing-in-memory; emerging memory; memory pooling and CXL fabrics; 3D DRAM post-2030 path

**Advanced Packaging and Substrate Materials** — The 2.5D/3D architecture continuum; TSMC CoWoS-S/L/R roadmap and capacity expansion; CoWoS-Photonics and CoWoP; SoIC, SoIC-X, SoIC-P hybrid-bonded stacks; Intel EMIB and Foveros; Samsung I-Cube/X-Cube/H-Cube; ABF supply oligopoly; glass-core substrates; interposer materials (silicon TSV, glass, organic RDL); hybrid bonding equipment ecosystem; HBM4 hybrid bonding adoption; OSAT capacity and Asian dominance

**Co-Packaged Optics and Silicon Photonics for AI** — The optical interconnect imperative; CPO architecture and two network layers; TSMC COUPE, CoWoS-Photonics, iOIS; CoWoP and the NVIDIA Rubin transition; ASE VIPack and the merchant photonics packaging layer; optical I/O chipelets (AyarLabs TeraPHY, Lightmatter Passage, Celestial AI / Marvell); switch silicon and co-packaged optical engines; silicon photonics foundries; photonics packaging materials supply chain; market sizing 2026–2036

**Thermal Management for AI Data Centers** — The thermal crisis at the package level; thermal interface materials (liquid metal TIM, solder TIM, diamond-based TIMs); heat spreaders, vapor chambers, heat pipes; cold plates and direct-to-chip liquid cooling; the cold plate supply chain bottleneck; single-phase and two-phase immersion cooling; PFAS challenge; microfluidic and in-package cooling; coolant distribution units, manifolds, and facility plumbing; market forecast 2024–2036

**Power Delivery and GaN/SiC Transition** — The power crisis from 12V to 48V to 800V HVDC; 48V tray architecture and OCP standard; 800V HVDC at the rack and the Rubin transition; SiC devices and substrate supply; GaN devices (lateral, vertical, cascode); GaN in AI server PSU applications; vertical GaN post-2027 trajectory; voltage regulator modules and multi-phase point-of-load; Monolithic Power Systems advantage in AI VRMs; package-integrated VRM; server PSUs and rack rectifier shelves; backside power delivery (Intel PowerVia, TSMC A16, Samsung BSPDN); market forecast 2024–2036

**Networking and Optical Materials** — The three network layers in an AI datacenter; switch silicon roadmap (Broadcom Tomahawk 6 Davison, NVIDIA Spectrum-X/Quantum-X); Ultra Ethernet Consortium; pluggable optical transceivers; volume transceiver suppliers; optical transceiver assembly (Fabrinet, Jabil, Luxshare); DSP and SerDes; Marvell's DSP business; Linear Pluggable Optics (LPO); III-V materials (InP, GaAs, GaN-Photonics); NICs, DPUs, and SmartNICs; cables, connectors, and DAC; market forecast 2024–2036

**Data Center Construction and Sustainability** — Power infrastructure (grid, on-site generation, SMRs); behind-the-meter natural-gas; nuclear restart and SMR procurement; renewable energy procurement at hyperscaler scale; switchgear and transformers; facility-level cooling architecture; construction supply chain and modular datacenter architecture; geographic concentration and site selection; PUE, WUE, and sustainability metrics; carbon-free energy accounting; embodied carbon; regulatory framework

**Edge GenAI Hardware** — AI smartphones and Apple Neural Engine evolution; AI PCs (NVIDIA, Snapdragon X Elite); NVIDIA Jetson and embedded AI; Jetson AGX Thor and humanoid robotics; automotive AI silicon (NVIDIA DRIVE Thor, Tesla FSD); humanoid robotics unit volumes and silicon revenue forecast; edge AI startup cohort; edge AI memory (LPDDR5X, on-chip SRAM, eMRAM); market forecast 2024–2036

**Regional Analysis** — Taiwan, South Korea, Japan, China, Southeast Asia and India, the United States, Europe and Israel; aggregate regional capture scenario analysis 2026–2036

**Supply Chain and Geopolitics** — The China strategy and sovereign stack; SMIC

and the EUV-free leading-edge path; CXMT and JHICC HBM ramp; US CHIPS Act implementation (TSMC Arizona, Samsung Taylor, Intel Foundry, Micron); European Chips Act; critical materials (rare earths, gallium and germanium, neon and specialty gases, specialty quartz and substrates); single-point-of-failure analysis; supply-chain resilience scenarios; sovereign AI as a strategic demand driver

Sustainability and Embodied Carbon — Operational emissions; cooling energy tax; embodied carbon in semiconductor manufacturing; PFC and process-gas problem; PFAS chemistry transition; renewable energy procurement; nuclear restart and SMR; heat recovery and district heating; circular economy; carbon accounting standards (Scope 1/2/3, EU CSRD, SEC); green manufacturing practices

Market Forecasts 2026–2036 — Total market Base case; Bull/Base/Bear scenarios; AI accelerator silicon, HBM, advanced packaging, photonics packaging, thermal, power, networking, datacenter construction, and edge AI sub-segment forecasts; regional capture forecast; customer-tier forecast; key forecast risks and sensitivities

Strategic Outlook — Five defining themes; choke-point map; strategic investment framework; M&A landscape and strategic consolidation through 2030; sensitivity analysis; strategic implications by stakeholder

Companies profiled include 1X Technologies, 3M, Acbel Polytech, Accelink Technologies, Achronix Semiconductor, Advanced Micro Devices (AMD), AGC (Asahi Glass), Agility Robotics, AheadComputing, Ajinomoto FineTechno (ABF), Akhan Semiconductor, Alibaba T-Head (PingTouGe), Alpha Assembly Solutions (MacDermid Alpha), Alphabet Inc. (Google), Amazon Web Services (AWS), Ambarella, Amber Semiconductor (AmberSemi), Amkor Technology, Amphenol Corporation, Anduril Industries, Apple Inc., Applied Materials, Appttronik, Arago, ASE Technology Holding (incl. SPIL), Asetek, Asia Vital Components (AVC), ASMPT, Asperitas, Astera Labs, Astrus, AT&S (Austria Technologie & Systemtechnik), Auras Technology, Avalanche Technology, Axelera AI, Axera Technology, AXT Inc., Ayar Labs, BE Semiconductor Industries (BESI), Biren Technology, Black Sesame Technologies, Blaize, Broadcom Inc., Cambricon Technologies, Cambridge GaN Devices (CGD), Carbice Corporation, Celero Communications, Cerebras Systems, Chemours Company, ChipAgents, Chipmind, ChipMOS Technologies, Chiral, Ciena, Cisco Systems, Claros, Coherent

Corp., ColorChip, Cooler Master Co., CoolIT Systems, CoreWeave Inc., Corintis, Corning Incorporated, Crossbar Inc., Crusoe Energy Systems, CXMT (ChangXin Memory Technologies), d-Matrix, DEEPX, Delta Electronics, DOW Inc., Dust Photonics, Eaton Corporation, EdgeCortex, EFFECT Photonics, Efficient Computer, Efficient Power Conversion (EPC), Element Six (e6), Eliyan, Empower Semiconductor, Engineered Fluids, Eoptolink Technology, Eridu, Etched.ai, Ethernovia, EuQlid, EV Group (EVG), Everspin Technologies, Fabric8Labs, Fabrinet, Femtum, Ferroelectric Memory Company (FMC), Figure AI, Fourier Intelligence, Foxconn Industrial Internet (FII), Foxconn Interconnect Technology (FIT), Frore Systems, FSP Group, Fujipoly, Furiosa AI, G42, Gaianixx, Galatek, Gigalight, Great Sky, Green Revolution Cooling (GRC), GreenWaves Technologies, Groq Inc., GS Microelectronics (GSME), Hailo Technologies, Henkel AG, Heraeus, Hesheng Silicon Industry, Hisense Broadband, Hitachi Energy, Hon Hai (Foxconn), Honeywell International, Horizon Robotics, Hua Tian Technology (HT-Tech), Huawei Technologies (HiSilicon), Hummink, Ibiden Co. Ltd., Iceotope Technologies, Iluvatar CoreX, Indium Corporation, Infineon Technologies AG, Innolight Technology, Innoscience Technology, Intel Corporation, Intel Foundry, IQE plc, JCET Group, JetCool Technologies, Kandou AI, Kaneka Corporation, Kinsus Interconnect Technology, Kioxia Holdings, Kneron, Kulicke & Soffa Industries (K&S), Kyocera Corporation, Lace Lithography, Lam Research, Lambda Inc., LG Innotek, Lightmatter, Liquid Wire Inc., LiquidStack, LiteOn Technology, LOTES Co., Lumentum Holdings, Lumotive, Luxshare Precision, M&I Materials, Macronix International, Maieutic Semiconductor, Majestic Labs, Marvell Technology, MatX, MediaTek, Mesh Optical Technologies, Meta Platforms, Microchip Technology, Micron Technology Inc., Microsoft Corporation, Mitsubishi Electric, Mobileye Global, Monolithic Power Systems (MPS), Montage Technology, Moore Threads Technology, Morphing Machines, Movandi, Multibeam Corporation, Murata Manufacturing, Mythic, Nan Ya PCB, Nanya Technology, Navitas Semiconductor, NcodiN, Neo Semiconductor, NeoGraf Solutions, NeoLogic, Netrasemi, NEURA Robotics, Neurophos, Normal Computing, NVIDIA Corporation, NXP Semiconductors, Olix, Omni Design Technologies, onsemi (ON Semiconductor), OpenLight, Optalysys, Opticore, Oracle Corporation (Oracle Cloud Infrastructure), Oxmiq Labs, Panasonic, Parker Chomerics, Patentix, Positron AI, Power Integrations, Powerchip Semiconductor (PSMC), PowerLattice, Powertech Technology, Primemas and more...

## Contents

### 1 EXECUTIVE SUMMARY

- 1.1 Key Findings
- 1.2 The Generative AI Hardware Bottleneck
- 1.3 Materials Value Chain at a Glance
- 1.4 Ten-Year Forecast Highlights
- 1.5 Strategic Implications for Asian Foundries, OSAT, Memory, Substrate, and Cooling Vendors
- 1.6 Differentiation vs. Adjacent Coverage
- 1.7 Major Market Players

### 2 THE COMPUTE STACK BEHIND GENERATIVE

- 2.1 Training vs. Inference Economics
  - 2.1.1 Pre-training, post-training, RLHF compute splits
  - 2.1.2 Inference token economics and serving infrastructure
  - 2.1.3 Test-time compute and reasoning-model demand
- 2.2 Cloud, Edge, and Sovereign AI
  - 2.2.1 Hyperscaler clusters at 100,000-GPU scale
  - 2.2.2 Enterprise on-prem and neocloud deployments
  - 2.2.3 Sovereign AI build-outs
  - 2.2.4 Edge inference cross-reference
- 2.3 Why Memory Bandwidth and Packaging Dominate Cost
  - 2.3.1 The memory wall in LLM serving
  - 2.3.2 HBM ASP as percentage of AI accelerator BOM
  - 2.3.3 CoWoS as the constraining bottleneck
- 2.4 Materials and Components as the New Bottleneck
- 2.5 Hyperscaler vs. Enterprise vs. Sovereign Capex
- 2.6 Company Profiles
  - 2.6.1 Alphabet Inc. (Google)
  - 2.6.2 Amazon Web Services (AWS)
  - 2.6.3 CoreWeave Inc.
  - 2.6.4 Crusoe Energy Systems
  - 2.6.5 G42
  - 2.6.6 Lambda Inc.
  - 2.6.7 Meta Platforms
  - 2.6.8 Microsoft Corporation

## 2.6.9 Oracle Corporation (Oracle Cloud Infrastructure)

### **3 AI ACCELERTOR SILICON**

#### 3.1 GPUs

3.1.1 NVIDIA roadmap: Hopper ? Blackwell ? Blackwell Ultra ? Rubin ? Rubin Ultra

3.1.2 NVL72 rack architecture and post-Rubin scale-up

3.1.3 AMD MI300X ? MI355X ? MI400 trajectory

3.1.4 Intel Gaudi and the post-Gaudi roadmap

#### 3.2 Custom Hyperscaler ASICs

3.2.1 Google TPU v5/v6/v7 and ML supercomputer architecture

3.2.2 AWS Trainium 2/3 and Inferentia

3.2.3 Microsoft Maia and Cobalt

3.2.4 Meta MTIA generations

3.2.5 ASIC NRE economics and break-even analysis

#### 3.3 Domain-Specific and Challenger Architectures

3.3.1 Cerebras WSE-3 wafer-scale

3.3.2 Groq LPU deterministic inference

3.3.3 SambaNova RDU and dataflow

3.3.4 Tenstorrent, d-Matrix, Etched, Rivos, Lightmatter

#### 3.4 Chinese AI Chip Ecosystem

3.4.1 Huawei Ascend 910C / 910D / 950

3.4.2 Cambricon, Biren, Moore Threads, Iluvatar CoreX

3.4.3 Alibaba T-Head Hanguang and PingTouGe

3.4.4 Domestic substitution timeline to gen-on-gen parity

#### 3.5 Process Nodes and Foundry Roadmaps

3.5.1 TSMC: N3 ? N3P ? N2 ? N2P ? A16 ? A14

3.5.2 Samsung Foundry: 3GAP ? 2GAP ? SF1.4

3.5.3 Intel Foundry: 18A ? 14A and external customer pipeline

3.5.4 SMIC: N+1 / N+2 and the EUV-free 5nm question

3.5.5 EUV and High-NA EUV adoption curves

#### 3.6 Wafer-Level Integration and Reticle Stitching

#### 3.7 Company Profiles 82 (53 company profiles)

### **4 AI-DRIVEN CHIP DESIGN (EDA)**

#### 4.1 The EDA Bottleneck in the AI Hardware Era

#### 4.2 The Recursive Loop: AI Designing AI Hardware

#### 4.3 The Incumbent EDA Vendors' AI Initiatives

- 4.4 The Startup Cohort: Four Distinct Approaches
  - 4.4.1 Agentic AI for digital design and verification
  - 4.4.2 Physics-AI for simulation and advanced packaging
  - 4.4.3 AI for analog and PCB design
  - 4.4.4 EDA-adjacent silicon and applied AI
- 4.5 Geographic Distribution
- 4.6 Market Forecast: AI-EDA Tools 2026–2036
- 4.7 Strategic Implications
- 4.8 Company profiles (6 company profiles)

## **5 HIGH BANDWIDTH MEMORY AND BEYOND**

- 5.1 HBM Architecture and TSV Stacking Fundamentals
- 5.2 HBM Generation Roadmap
  - 5.2.1 HBM3 / HBM3E specifications and deployment
  - 5.2.2 HBM4 / HBM4E: pin width doubling and base-die logic
  - 5.2.3 HBM5 / HBM5E: 2031–2036 architecture directions
- 5.3 Memory Makers and Capacity Outlook
  - 5.3.1 SK hynix strategy, products, capex through 2030
  - 5.3.2 Samsung HBM3E re-qualification and HBM4 catch-up
  - 5.3.3 Micron HBM3E entry and AI customer share gains
  - 5.3.4 HBM bit-shipment and wafer-capacity forecasts
- 5.4 Custom HBM (cHBM) and Base-Die Innovation
  - 5.4.1 Customer-specific HBM with NVIDIA, Broadcom, Google
  - 5.4.2 Standard vs custom HBM revenue split through 2030
- 5.5 Compute-in-Memory and Processing-in-Memory at Scale
- 5.6 Emerging Memory for AI Datacenters
  - 5.6.1 Storage-class memory after 3D XPoint
- 5.7 Memory Pooling and CXL Fabrics
- 5.8 3D DRAM — The Post-2030 Path
- 5.9 Company Profiles 143 (23 company profiles)

## **6 ADVANCED PACKAGING AND SUBSTRATE MATERIALS**

- 6.1 The 2.5D / 3D Architecture Continuum
- 6.2 TSMC CoWoS and the Capacity Constraint
  - 6.2.1 CoWoS-S, CoWoS-L, CoWoS-R roadmap
  - 6.2.2 CoWoS-Photonics and CoWoP
  - 6.2.3 CoWoS capacity expansion: 2024 vs. 2026 vs. 2028 vs. 2030

- 6.2.4 SoIC, SoIC-X, SoIC-P: Hybrid-Bonded Stacks
- 6.3 Intel and Samsung Advanced Packaging
  - 6.3.1 Intel: EMIB, EMIB-T, Foveros, Foveros Direct, Foveros Omni
  - 6.3.2 Samsung: I-Cube, X-Cube, H-Cube
- 6.4 Substrate Technologies (ABF, FC-BGA)
  - 6.4.1 ABF supply oligopoly
  - 6.4.2 Glass core substrate (Intel, ASE, SCHOTT)
- 6.5 Interposer Materials (Silicon TSV, Glass, Organic RDL)
- 6.6 Hybrid Bonding and Copper-to-Copper Interconnect
  - 6.6.1 Hybrid bonding equipment ecosystem
  - 6.6.2 HBM4 adoption of hybrid bonding
- 6.7 OSAT Capacity and Asian Dominance
- 6.8 Advanced Packaging Materials Suppliers
- 6.9 Company Profiles 170 (56 company profiles)

## **7 CO-PACKAGED OPTICS AND SILICON PHOTONICS FOR AI**

- 7.1 The Optical Interconnect Imperative
- 7.2 CPO Architecture and the Two Network Layers
- 7.3 TSMC COUPE, CoWoS-Photonics, iOIS
  - 7.3.1 TSMC photonics design ecosystem
  - 7.3.2 CoWoP and the NVIDIA Rubin transition
- 7.4 ASE VIPack and the Merchant Photonics Packaging Layer
- 7.5 Optical I/O Chiplets: AyarLabs, Lightmatter, Celestial AI
  - 7.5.1 AyarLabs TeraPHY
  - 7.5.2 Lightmatter Passage
  - 7.5.3 Celestial AI Photonic Fabric and the Marvell acquisition
- 7.6 Switch Silicon and Co-Packaged Optical Engines
- 7.7 Silicon Photonics Foundries
- 7.8 Photonics Packaging Materials and Supply Chain
- 7.9 Market Sizing for Photonics Packaging 2026–2036
- 7.10 Company Profiles 214 (28 company profiles)

## **8 THERMAL MANAGEMENT FOR AI DATA CENTERS**

- 8.1 The Thermal Crisis: Power Density at the Package Level
- 8.2 Thermal Interface Materials (TIMs)
  - 8.2.1 Liquid metal TIM and the gallium corrosion problem
  - 8.2.2 Solder TIM (indium and SnAg)

- 8.2.3 Diamond-based TIMs and emerging materials
- 8.3 Heat Spreaders, Vapor Chambers, and Heat Pipes
- 8.4 Cold Plates and Direct-to-Chip Liquid Cooling
  - 8.4.1 Cold plate design and microchannel geometry
  - 8.4.2 The cold plate supply chain bottleneck
- 8.5 Immersion Cooling
  - 8.5.1 Single-phase immersion: mineral oil and synthetic dielectrics
  - 8.5.2 Two-phase immersion: fluorocarbons and the PFAS challenge
- 8.6 Microfluidic and In-Package Cooling
  - 8.6.1 Microfluidic ecosystem and the first commercial applications
  - 8.6.2 Coolant Distribution Units, Manifolds, and Facility Plumbing
- 8.7 Market Forecast: AI-Tied Thermal Management 2024–2036
- 8.8 Company Profiles 242 (40 company profiles)

## **9 POWER DELIVERY AND GAN/SIC TRANSITION**

- 9.1 The Power Crisis: From 12V to 48V to 800V HVDC
- 9.2 The Power Hierarchy: System ? Board ? Package ? Die
  - 9.2.1 48V tray architecture and the OCP standard
  - 9.2.2 800V HVDC at the rack and the Rubín transition
- 9.3 SiC Devices and Substrate Supply
  - 9.3.1 SiC substrate supply: the bottleneck
- 9.4 GaN Devices: Lateral, Vertical, Cascode
  - 9.4.1 GaN switching speed and AI server PSU applications
  - 9.4.2 Vertical GaN: the post-2027 trajectory
- 9.5 Voltage Regulator Modules and Multi-Phase Point-of-Load
  - 9.5.1 The Monolithic Power Systems advantage in AI VRMs
  - 9.5.2 Vertical power delivery and the package-integrated VRM
- 9.6 Server Power Supply Units and Rack Rectifier Shelves
- 9.7 Backside Power Delivery (BSPDN)
  - 9.7.1 Intel PowerVia (18A)
  - 9.7.2 TSMC backside power (A16)
  - 9.7.3 Samsung BSPDN
- 9.8 Market Forecast: AI Datacenter Power Semiconductors 2024–2036
- 9.9 Company Profiles 277 (42 company profiles)

## **10 NETWORKING AND OPTICAL MATERIALS**

- 10.1 The Three Network Layers in an AI Datacenter

- 10.2 Switch Silicon Roadmap
  - 10.2.1 Tomahawk 6 Davison and the CPO inflection
  - 10.2.2 NVIDIA Spectrum-X and Quantum-X
  - 10.2.3 Ultra Ethernet Consortium (UEC)
- 10.3 Pluggable Optical Transceivers
  - 10.3.1 Volume optical transceiver suppliers
  - 10.3.2 Optical transceiver assembly: Fabrinet, Jabil, Luxshare
- 10.4 DSP and SerDes for Optical Transceivers
  - 10.4.1 Marvell's DSP business and the AI optical transceiver
  - 10.4.2 Linear Pluggable Optics (LPO) and the DSP-less transceiver
- 10.5 III-V Materials Layer: InP, GaAs, GaN-Photonics
- 10.6 NICs, DPUs, and SmartNICs
- 10.7 Cables, Connectors, and Direct Attach Copper
- 10.8 Market Forecast: AI-Tied Networking and Optical 2024–2036
- 10.9 Company Profiles 310 (36 company profiles)

## **11 DATA CENTER CONSTRUCTION AND SUSTAINABILITY**

- 11.1 The AI Datacenter Buildout: Scale and Scope
- 11.2 Power Infrastructure: Grid, On-Site Generation, and SMRs
  - 11.2.1 Behind-the-meter natural-gas generation
  - 11.2.2 Nuclear restart and Small Modular Reactor procurement
  - 11.2.3 Renewable energy procurement at hyperscaler scale
  - 11.2.4 Switchgear and transformers: the silent bottleneck
- 11.3 Facility-Level Cooling Architecture
- 11.4 Construction Supply Chain and Modular Datacenter Architecture
- 11.5 Geographic Concentration and Site Selection
  - 11.5.1 The Top 12 AI Datacenter Regions (2026)
  - 11.5.2 Climate as a constraint
- 11.6 PUE, WUE, and Sustainability Metrics
  - 11.6.1 Carbon-Free Energy (CFE) accounting
  - 11.6.2 Embodied carbon and circular economy
- 11.7 Regulatory Framework
  - 11.7.1 Permit and interconnection timelines
- 11.8 Market Forecast: AI Datacenter Construction Supply Chain 2024–2036

## **12 EDGE GEN AI HARDWARE**

- 12.1 The Edge AI Taxonomy

## 12.2 AI Smartphones

### 12.2.1 Apple Neural Engine evolution

## 12.3 AI PCs

### 12.3.1 NVIDIA's AI PC entry

### 12.3.2 Snapdragon X Elite and Qualcomm's PC push

## 12.4 NVIDIA Jetson and the Embedded AI Platform

### 12.4.1 Jetson AGX Thor and humanoid robotics

## 12.5 Automotive AI Silicon

### 12.5.1 NVIDIA DRIVE Thor and the L4 autonomous driving platform

### 12.5.2 Tesla FSD and the captive silicon path

## 12.6 Humanoid Robotics: The Emerging Edge AI Compute Frontier

### 12.6.1 Humanoid robot unit volumes and silicon revenue forecast

## 12.7 Edge AI Accelerator Start-ups

## 12.8 Edge AI Memory: LPDDR5X, On-Chip SRAM, eMRAM

## 12.9 Market Forecast: Edge AI Silicon 2024–2036

## 12.10 Company Profiles 349 (51 company profiles)

# **13 REGIONAL ANALYSIS: GEOGRAPHY OF THE GENAI HARDWARE SUPPLY CHAIN**

## 13.1 The Asian Concentration

### 13.2 Taiwan

#### 13.2.1 The TSMC scale

#### 13.2.2 The Taiwan supply chain depth

#### 13.2.3 Taiwan's geographic concentration risk

### 13.3 South Korea

#### 13.3.1 SK hynix as the strategic anchor

#### 13.3.2 Samsung: vertical integration across the stack

#### 13.3.3 Korean specialty positions

### 13.4 Japan

#### 13.4.1 Kumamoto and the broader Japanese fab expansion

### 13.5 China

#### 13.5.1 Chinese domestic AI silicon volume and trajectory

#### 13.5.2 The SMIC constraint

#### 13.5.3 China's strength layers

### 13.6 Southeast Asia and India

#### 13.6.1 Malaysian AI infrastructure

#### 13.6.2 India's emerging fab and OSAT capacity

#### 13.6.3 ASEAN AI cloud and sovereign-AI initiatives

## 13.7 The United States

### 13.7.1 The CHIPS Act build-out

### 13.7.2 The US labour and supply chain constraints

## 13.8 Europe and Israel

### 13.8.1 ASML

### 13.8.2 European Chips Act and the limits of European industrial policy

### 13.8.3 Israel's specialty position

## 13.9 The Rest of World: Niche Capabilities and Sovereign Ambitions

## 13.10 Aggregate Regional Capture: Scenario Analysis 2026–2036

# 14 SUPPLY CHAIN AND GEOPOLITICS

## 14.1 The Defining Tensions

## 14.2 The China Strategy: Sovereign Stack and Domestic Substitution

### 14.2.1 SMIC's role and the EUV-free leading-edge path

### 14.2.2 The CXMT and JHICC HBM ramp

### 14.2.3 China's wafer-fab equipment indigenisation

## 14.3 US CHIPS Act Implementation and Domestic Reshoring

### 14.3.1 TSMC Arizona

### 14.3.2 Samsung Taylor

### 14.3.3 Intel Foundry

### 14.3.4 Micron's CHIPS-supported expansion

### 14.3.5 The labour and ecosystem constraints

## 14.4 European Chips Act and Strategic Autonomy

### 14.4.1 The European specialty position

## 14.5 The Critical Materials Layer

### 14.5.1 Rare earths

### 14.5.2 Gallium and germanium

### 14.5.3 Neon and specialty gases

### 14.5.4 Specialty quartz, silicon, and substrates

## 14.6 Single-Point-of-Failure Analysis

## 14.7 Scenarios for Supply Chain Resilience

### 14.7.1 The "successful diversification" scenario (Bull case for resilience)

### 14.7.2 The "concentrated capacity" scenario (Base case)

### 14.7.3 The "geopolitical disruption" scenario (Bear case for resilience)

## 14.8 Sovereign AI as a Strategic Demand Driver

# 15 SUSTAINABILITY AND EMBODIED CARBON

- 15.1 The Sustainability Stakes
- 15.2 Operational Emissions: Training, Inference, and the Cooling Energy Tax
  - 15.2.1 Training versus inference: the dominant share
- 15.3 Embodied Carbon in Semiconductor Manufacturing
  - 15.3.1 The PFC and process-gas problem
  - 15.3.2 Embodied carbon at the device level
  - 15.3.3 Server-level and facility-level embodied carbon
- 15.4 Water, Chemicals, and Resource Intensity
  - 15.4.1 PFAS chemistry and the transition
- 15.5 Renewable Energy Procurement at Hyperscaler Scale
  - 15.5.1 Nuclear restart and SMR as carbon-free baseload
  - 15.5.2 On-site natural gas: the carbon offset
- 15.6 Heat Recovery, Circular Economy, and End-of-Life
  - 15.6.1 Heat recovery and district heating
  - 15.6.2 Circular economy and component reuse
- 15.7 Carbon Accounting Standards and Corporate Disclosure
  - 15.7.1 Scope 1, 2, 3 framework
  - 15.7.2 EU Corporate Sustainability Reporting Directive
  - 15.7.3 SEC climate disclosure rules
  - 15.7.4 Carbon pricing and offsets
- 15.8 Green Manufacturing Practices at Major Suppliers
  - 15.8.1 Process gas abatement
  - 15.8.2 Water recycling and reuse
- 15.9 Market and Regulatory Outlook 2026–2036
  - 15.9.1 Carbon-related regulatory tightening
  - 15.9.2 Embodied-carbon-conscious procurement
  - 15.9.3 The carbon-aware AI compute frontier

## **16 MARKET FORECASTS: GEN AI HARDWARE 2026-2036**

- 16.1 Forecast Methodology and Framework
- 16.2 Total GenAI Hardware Market — Base Case Forecast
- 16.3 Bull/Base/Bear Scenarios at Aggregate Level
- 16.4 AI Accelerator Silicon Sub-Segment Forecast
  - 16.4.1 Merchant vs. captive ASIC share trajectory
  - 16.4.2 China sovereign-stack AI silicon trajectory
- 16.5 HBM and Memory Sub-Segment Forecast
- 16.6 Advanced Packaging Sub-Segment Forecast
- 16.7 Photonics Packaging Sub-Segment Forecast

- 16.8 Thermal Management Sub-Segment Forecast
- 16.9 Power Delivery Sub-Segment Forecast
- 16.10 Networking and Optical Sub-Segment Forecast
- 16.11 Datacenter Construction Supply Chain Sub-Segment Forecast
- 16.12 Edge AI Silicon Sub-Segment Forecast
- 16.13 Regional Capture Forecast
- 16.14 Customer Tier Forecast
- 16.15 Key Forecast Risks and Sensitivities
  - 16.15.1 The CapEx normalisation risk
  - 16.15.2 The Taiwan concentration risk
  - 16.15.3 Model training economics
  - 16.15.4 Chinese sovereign-stack acceleration
  - 16.15.5 Power infrastructure constraints

## **17 STRATEGIC OUTLOOK**

- 17.1 The Five Defining Themes of the GenAI Hardware Decade
- 17.2 The Choke-Point Map
- 17.3 The Strategic Investment Framework
- 17.4 M&A Landscape and Strategic Consolidation
  - 17.4.1 Photonics consolidation
  - 17.4.2 Memory and HBM consolidation
  - 17.4.3 Equipment and tools consolidation
  - 17.4.4 AI silicon start-up consolidation
  - 17.4.5 Forward M&A trajectory through 2030
- 17.5 Sensitivity Analysis
- 17.6 Strategic Implications by Stakeholder
  - 17.6.1 For AI accelerator silicon designers
  - 17.6.2 For hyperscalers and AI cloud operators
  - 17.6.3 For memory manufacturers
  - 17.6.4 For foundries
  - 17.6.5 For OSATs and substrate suppliers
  - 17.6.6 For thermal and power infrastructure suppliers
  - 17.6.7 For photonics packaging participants
  - 17.6.8 For governments and policymakers
- 17.7 What Could Change This Forecast
  - 17.7.1 Upside surprises
  - 17.7.2 Downside surprises
  - 17.7.3 Structural rather than cyclical risk

## **18 APPENDIX**

### 18.1 Forecast Methodology

- 18.1.1 Unit volume forecast construction
- 18.1.2 ASP and content-per-unit forecast construction
- 18.1.3 Scenario construction
- 18.1.4 Cross-validation

### 18.2 Definitions and Terminology

- 18.2.1 AI accelerator silicon categories
- 18.2.2 Memory technology categories
- 18.2.3 Packaging terminology
- 18.2.4 Photonics terminology
- 18.2.5 Thermal terminology
- 18.2.6 Power terminology
- 18.2.7 Networking terminology
- 18.2.8 Geographic and customer terminology

### 18.3 Abbreviations

### 18.4 Sources and References

- 18.4.1 Primary research
- 18.4.2 Company financial disclosures
- 18.4.3 Industry-association and government statistics
- 18.4.4 Cross-reference industry reports
- 18.4.5 Technical and scientific literature

### 18.5 Forecast Scope, Limitations, and Disclaimers

- 18.5.1 Forecast scope
- 18.5.2 Forecast limitations
- 18.5.3 Disclaimers

### 18.6 Detailed Year-by-Year Forecast Outputs

## List Of Tables

### LIST OF TABLES

- Table 1. Headline Findings Summary (Base Case)
- Table 2. Ten-Year Forecast Summary: GenAI Hardware Materials Market 2026–2036 (US \$B, Base Case)
- Table 3. Top Ten Strategic Conclusions Mapped to Stakeholder Type
- Table 4. Training vs. Inference Hardware Mix Comparison
- Table 5. Silicon Content per 100 MW AI Training Facility (Reference BoM)
- Table 6. Cost-per-Token by Model Size and Hardware Configuration 2024–2040 (USD per million output tokens)
- Table 7. Sovereign AI Build-Outs by Country 2025–2030
- Table 8. AI Accelerator Memory Requirements 2024–2030F
- Table 9. US and Chinese Hyperscaler Capex Summary 2021–2026 (US \$B)
- Table 10. GPU Specifications: NVIDIA Blackwell, Rubin; AMD MI350X, MI450 (2024–2026)
- Table 11. Rack-Scale GPU Platform Comparison
- Table 12. AI ASIC Specifications: Google, AWS, Microsoft, Meta (2024–2026)
- Table 13. AI ASIC Technology Specification Database (All Major Vendors)
- Table 14. Chinese Data Center Processor Manufacturer Overview
- Table 15. China AI Chip Capability Gap Assessment by Workload Type
- Table 16. Semiconductor Process Node Roadmap 2024–2030
- Table 17. TSMC Node Roadmap: N3, N2, A16, A14 Specs and Timeline
- Table 18. Wafer-Scale Accelerator Yield Economics: Cerebras WSE-3 and Tesla Dojo
- Table 19. Incumbent EDA Vendor AI Initiatives vs. Startup Cohort
- Table 20. AI-EDA Approaches by Design-Flow Stage
- Table 21. AI-EDA Market Forecast 2026–2036
- Table 22. HBM Generation Technical Specifications HBM2E to HBM5
- Table 23. HBM Bonding Integration Roadmap and Vendor Mapping
- Table 24. HBM Market Share by Supplier 2022–2028F (%)
- Table 25. HBM Customer Demand Breakdown: NVIDIA, Google, AMD, Hyperscalers 2024–2028F
- Table 26. Custom HBM Players, Products, Design Roadmaps
- Table 27. Standard vs. Custom HBM Revenue Forecast 2024–2030F (US \$M)
- Table 28. Near-Memory and In-Memory Computing Landscape
- Table 29. Resistive Non-Volatile Memory Technologies
- Table 30. Storage-Class Memory Technology Comparison
- Table 31. CXL Switch Silicon Vendors and Capability Matrix

- Table 32. 3D DRAM Technology Readiness Assessment by Player 2026
- Table 33. Advanced Packaging Technology Comparison: 2.5D and 3D Options
- Table 34. CoWoS Capacity Forecast by Sub-Variant 2024–2036 (k wafers/month equivalent)
- Table 35. TSMC SoIC Variants: Specifications and AI Customer Adoption
- Table 36. Comparative Advanced Packaging Roadmap: TSMC vs. Intel vs. Samsung
- Table 37. Substrate Suppliers for AI Accelerator Packages
- Table 38. Substrate Demand Forecast for AI Packages 2024–2036 (k units/month)
- Table 39. Interposer Material Comparison: Silicon TSV vs. Glass vs. Organic RDL
- Table 40. Hybrid Bonding Adoption Roadmap for DRAM Applications 2023–2030
- Table 41. OSAT Capacity and Revenue Concentration 2024–2030
- Table 42. Advanced Packaging Materials Suppliers
- Table 43. Migration Trajectory from Copper to Optical Across the Two Network Layers
- Table 44. Key Technology Building Blocks for Co-Packaged Optics
- Table 45. TSMC Photonics Packaging Capabilities
- Table 46. Merchant Photonics Packaging Platform Comparison
- Table 47. Optical I/O Chiplet Vendor Comparison
- Table 48. AI-Switch Silicon Roadmap with CPO Integration
- Table 49. Silicon Photonics Foundry Capability Matrix
- Table 50. CPO Supply Chain Critical Materials and Suppliers
- Table 51. Photonics Packaging Revenue Forecast for AI Applications 2024–2036 (US \$B)
- Table 52. Cooling Technologies for High-Performance AI Processors
- Table 53. Thermal Interface Material Categories and Suppliers
- Table 54. TIM Properties for AI Accelerator Applications
- Table 55. TIM Revenue Forecast for AI Datacenter Applications 2024–2036 (US \$M)
- Table 56. Heat Spreader and Vapor Chamber Suppliers
- Table 57. Heat Spreader and Heat Sink Revenue Forecast 2024–2036 (US \$M)
- Table 58. Cold Plate Suppliers for AI Servers
- Table 59. Liquid Cooling Adoption Share in New AI Datacenter Deployments
- Table 60. Immersion Cooling Fluid Categories and Suppliers
- Table 61. Immersion Cooling System Suppliers
- Table 62. Microfluidic Cooling Technology Comparison
- Table 63. Facility Liquid Cooling Infrastructure Suppliers
- Table 64. AI-Tied Thermal Management Revenue Forecast 2024–2036 (US \$B)
- Table 65. Power Delivery Hierarchy in AI Servers
- Table 66. Comparison of 48V and 800V HVDC Rack Architectures
- Table 67. SiC vs. GaN vs. Silicon Power Device Comparison
- Table 68. SiC Substrate and Device Suppliers

- Table 69. GaN Device Manufacturers and Application Focus
- Table 70. AI VRM Controller and Power Stage Suppliers
- Table 71. Server Power Supply Unit Suppliers
- Table 72. Backside Power Delivery Adoption Roadmap
- Table 73. AI Datacenter Power Semiconductor Revenue Forecast 2024–2036 (US \$B)
- Table 74. The Three Networking Layers in an AI Datacenter
- Table 75. AI Switch Silicon Roadmap
- Table 76. Optical Transceiver Form Factor and Data Rate Roadmap
- Table 77. Optical Transceiver Module Suppliers for AI Datacenters
- Table 78. Optical DSP Suppliers and Application Mapping
- Table 79. III-V Substrate Materials Suppliers for AI Optical Transceivers
- Table 80. NIC, DPU, and SmartNIC Suppliers
- Table 81. Cable, Connector, and Fiber Suppliers for AI Datacenters
- Table 82. AI-Tied Networking and Optical Revenue Forecast 2024–2036 (US \$B)
- Table 83. AI Datacenter CAPEX Breakdown (100 MW Training Facility, 2026 Reference)
- Table 84. Hyperscaler Power Procurement Strategies (2025 Snapshot)
- Table 85. Major Switchgear, Transformer, and Power Infrastructure Suppliers
- Table 86. Facility Cooling Infrastructure Suppliers
- Table 87. Major AI Datacenter Construction Companies and Operators
- Table 88. Construction Engineering and EPC Firms with Major AI Datacenter Practice
- Table 89. PUE Targets and Achievement at Major Hyperscalers (2025)
- Table 90. AI-Tied Datacenter Construction Supply Chain Revenue Forecast 2024–2036 (US \$B)
- Table 91. Edge AI NPU Performance by Application Segment
- Table 92. Flagship Smartphone AI Processor Comparison (2026)
- Table 93. Evolution of Apple Neural Engine AI Performance (2017–2026)
- Table 94. AI PC Silicon Platform Comparison (2026)
- Table 95. AI PC On-Device LLM Inference Capability (2026)
- Table 96. NVIDIA Jetson Product Line (2026)
- Table 97. Automotive AI Silicon Platforms (2026)
- Table 98. Humanoid Robot Compute Platforms (2026)
- Table 99. Edge AI Start-up Landscape
- Table 100. Edge AI Memory Suppliers and Categories
- Table 101. Edge AI Silicon Revenue Forecast 2024–2036 (US \$B)
- Table 102. Regional Capture of GenAI Hardware Bill of Materials, 2026 Base Case
- Table 103. Taiwan AI Hardware Supply Chain by Capability Layer
- Table 104. Korea AI Hardware Supply Chain by Capability Layer
- Table 105. Japan AI Hardware Supply Chain by Capability Layer

- Table 106. China AI Hardware Supply Chain by Capability Layer
- Table 107. Southeast Asia and India AI Hardware Supply Chain
- Table 108. United States AI Hardware Supply Chain by Capability Layer
- Table 109. Europe and Israel AI Hardware Supply Chain
- Table 110. Regional GenAI Hardware BoM Capture by Scenario (% of Global BoM Value)
- Table 111. Major US Export Control Actions Affecting AI Hardware (2019–2026)
- Table 112. Chinese Wafer-Fab Equipment Companies and Capability Status
- Table 113. Major CHIPS Act-Funded Semiconductor Projects
- Table 114. Critical Materials Supply Chain Concentration for AI Hardware
- Table 115. Top Single-Point-of-Failure Risks in the GenAI Hardware Supply Chain
- Table 116. Supply Chain Diversification Scenario Outcomes 2030
- Table 117. Lifecycle Carbon Footprint by AI Chip Type
- Table 118. AI Carbon Footprint Examples and Mitigation Strategies
- Table 119. Estimated Embodied Carbon Across the AI Hardware Hierarchy
- Table 120. Water Consumption Profile for AI Hardware Manufacturing and Operations
- Table 121. Hyperscaler Renewable Energy and Nuclear Procurement (2025 Snapshot)
- Table 122. Lifecycle and End-of-Life Treatment for AI Hardware
- Table 123. Major Corporate Carbon Commitments Affecting AI Hardware Procurement
- Table 124. Green Manufacturing Initiatives by Major Semiconductor Suppliers
- Table 125. Forecast Methodology and Key Assumptions
- Table 126. Total GenAI Hardware Market by Major Segment, Base Case (US \$B)
- Table 127. GenAI Hardware Aggregate Market Across Three Scenarios, 2026–2036 (US \$B, excl. construction supply chain)
- Table 128. AI Accelerator Silicon Sub-Segment Forecast 2024–2036 (US \$B)
- Table 129. HBM and AI-Tied Memory Sub-Segment Forecast 2024–2036 (US \$B)
- Table 130. Advanced Packaging Sub-Segment Forecast 2024–2036 (US \$B, AI-tied)
- Table 131. Photonics Packaging Sub-Segment Forecast 2024–2036 (US \$B)
- Table 132. Thermal Management Sub-Segment Forecast 2024–2036 (US \$B, AI-tied)
- Table 133. Power Delivery (AI Datacenter Tied) Sub-Segment Forecast 2024–2036 (US \$B)
- Table 134. Networking and Optical (AI-Tied) Sub-Segment Forecast 2024–2036 (US \$B)
- Table 135. Datacenter Construction Supply Chain Sub-Segment Forecast 2024–2036 (US \$B)
- Table 136. Edge AI Silicon Sub-Segment Forecast 2024–2036 (US \$B)
- Table 137. Regional GenAI Hardware BoM Capture Forecast, 2026–2036, Base Case (%)
- Table 138. Total GenAI Hardware Demand by Customer Tier, Base Case 2026–2036

(US \$B, excl. construction supply chain)

Table 139. The Five Defining Themes: Strategic Implications by Layer

Table 140. The Top 15 Strategic Choke Points in the GenAI Hardware Supply Chain

Table 141. Strategic Tier Classification of GenAI Hardware Sub-Segments

Table 142. Notable GenAI Hardware M&A and Strategic Investments 2020–2026

Table 143. Sensitivity of Base Case 2030 Forecast to Key Assumptions

Table 144. Detailed Year-by-Year Total Forecast, Base Case (US \$B, excl. DC construction supply chain)

Table 145. Detailed Year-by-Year Total Forecast Across All Three Scenarios (US \$B, excl. DC construction supply chain)

## List Of Figures

### LIST OF FIGURES

- Figure 1. Five Compute-Scaling Walls and Their Material Solutions
- Figure 2. Generative AI Hardware Materials Value-Chain Layer Map
- Figure 3. Base-Case Forecast Stacked-Area Visualisation 2026–2036
- Figure 4. Bull, Base, and Bear Scenario Comparison 2026–2036
- Figure 5. Asia-Pacific Capture Rate of GenAI Hardware Value 2026–2036
- Figure 6. AI Data Centre Silicon Content Map
- Figure 7. Inference Token Economics by Model Size
- Figure 8. Sovereign AI Capex Pipeline 2024–2030 by Geography
- Figure 9. Generative AI Compute Demand Scaling vs. Electrical Interconnect Capacity
- Figure 10. AI Accelerator BoM Decomposition: Where the Dollars Go
- Figure 11. Annual GenAI-Driven AI Hardware Demand Pool 2024–2030
- Figure 12. NVIDIA GPU Architecture Evolution: Volta to Post-Blackwell Timeline
- Figure 13. Rack-Scale GPU Architecture: NVL72 and Next-Generation Platforms
- Figure 14. Hyperscaler ASIC Roadmap Comparison
- Figure 15. Hyperscaler ASIC vs. Merchant GPU Share of Datacenter AI Compute 2024–2036
- Figure 16. AI ASIC Start-Up Landscape by Funding Stage
- Figure 17. GPU vs. AI ASIC Performance per Watt Comparison 2022–2026
- Figure 18. China Semiconductor Capability Map: Node vs. Supply-Chain Layer
- Figure 19. China AI Chip Roadmap vs. NVIDIA / AMD: Parity Distance by Generation
- Figure 20. Leading-Edge Foundry Roadmap Comparison 2023–2036 (Gantt)
- Figure 21. HBM Architecture: Die-Stack Cross-Section
- Figure 22. HBM Bandwidth Evolution HBM1 to HBM5
- Figure 23. HBM4 Die-to-Wafer Bonding Integration Scheme
- Figure 24. HBM Market Share by Supplier 2022–2028F
- Figure 25. SK hynix HBM Strategy and Roadmap
- Figure 26. Samsung HBM Strategy and Roadmap
- Figure 27. Micron HBM Strategy and Roadmap
- Figure 28. HBM Customer Demand Breakdown by AI Accelerator
- Figure 29. Custom HBM Architecture: Co-Design Concept
- Figure 30. Custom HBM Share of Total HBM Bit Demand 2026–2036
- Figure 31. Near-Memory vs. PIM Architecture Comparison
- Figure 32. CXL Memory Pooling Architecture and Vendor Map
- Figure 33. 3D DRAM Concept Architectures
- Figure 34. Monolithic Die vs. Chiplet Architecture: Yield and Cost

- Figure 35. Chiplet Interconnect Technology Spectrum
- Figure 36. CoWoS Integration: GPU + HBM on Silicon Interposer
- Figure 37. CoWoS Capacity Expansion Roadmap
- Figure 38. OSAT Revenue Concentration by Geography 2024–2036
- Figure 39. Compute Demand vs. Interconnect Bandwidth Gap
- Figure 40. Photonics Packaging Revenue Forecast for AI Applications 2024–2036
- Figure 41. AI Accelerator TDP and Cooling Architecture Trajectory 2022–2036
- Figure 42. Liquid Cooling Adoption Trajectory in AI Datacenter Deployments
- Figure 43. Power Density at AI Server Rack: From 30 kW to 600 kW per Rack
- Figure 44. Wide-Bandgap Power Semiconductor Material Properties Comparison
- Figure 45. Edge AI Performance and Power Envelope Map
- Figure 46. Total GenAI Hardware Market 2024–2036 by Segment, Base Case
- Figure 47. GenAI Hardware Market Bull/Base/Bear Scenarios 2024–2036
- Figure 48. Sensitivity of 2030 Forecast to Key Variables

## I would like to order

Product name: Generative AI Hardware Materials Market 2026-2036-Semiconductors, Memory, Packaging, and Thermal Management

Product link: <https://marketpublishers.com/r/GA36C20C2884EN.html>

Price: US\$ 1,600.00 (Single User License / Electronic Delivery)

If you want to order Corporate License or Hard Copy, please, contact our Customer Service:

[info@marketpublishers.com](mailto:info@marketpublishers.com)

## Payment

To pay by Credit Card (Visa, MasterCard, American Express, PayPal), please, click button on product page <https://marketpublishers.com/r/GA36C20C2884EN.html>