

Chapter 16: The Economics of AI Chips

chipsupplychain.org/component-costs/

CITE AS Sharma, Yashvardhan. "The Economics of AI Chips." *Chip Supply Chain: How AI Chips Are Made*. Edition of 20 October 2026. <https://chipsupplychain.org/component-costs/>

CHAPTER

16

THE ECONOMICS

The Economics of AI Chips

Epoch AI estimated in December 2025 that Nvidia's B200 AI chip cost about \$6,400 to build, mostly for memory and packaging. Its processor silicon cost about \$900.

IN PLAIN TERMS

A bill of materials is the list of parts inside a product and what each one costs. Epoch AI estimated in December 2025 that Nvidia's B200, an AI chip from its Blackwell line, cost about \$6,400 to build ¹. The two logic dies, the pieces of silicon that do the calculating, cost about \$900 of that ¹. The stacked memory beside them accounted for about half, and the packaging that wires the dies and memory together cost more than the dies ¹. Three firms, SK hynix, Samsung and Micron, dominate the world market for that memory ³. One Taiwanese firm, TSMC, does that packaging ¹.

¹ NVIDIA's B200 costs around \$6,400 to produce. Epoch AI · 10 December 2025 [A] epoch.ai/data-insights/b200-cost-breakdown

³ Understanding the Biden Administration's Updated Export Controls. Center for Strategic and International Studies · 11 December 2024 [A] www.csis.org/analysis/understanding-biden-administrations-updated-export-controls

IN SHORT

Reports in March 2024 put the B200's selling price at \$30,000 to \$40,000 per chip ¹. Epoch AI notes that most Blackwell revenue comes from servers and rack-scale systems, which may carry lower margins ¹. It says Nvidia's Blackwell margin may therefore be lower than the gap between cost and price suggests ¹.

The estimate's memory price of \$14 to \$17 per gigabyte predates the 2026 rise in memory makers' sales and margins ⁵. Micron's gross margin, the share of sales left after production costs, rose from 39.8 percent for the year to August 2025 to 86.8 percent in the quarter to September 2026 ¹¹. Nvidia's was 75.0 percent in the quarter to July 2026 ⁷.

In December 2025 Epoch AI estimated that building one of Nvidia's Blackwell B200 chips cost about \$6,400 ¹. Reports in March 2024 put its selling price at \$30,000 to \$40,000 per chip ¹. Epoch AI's chip sales model, with data through October 2025, puts Nvidia's compute revenue per GPU in its GB200 rack systems at \$33,000 to \$42,000 ². That figure leaves out the networking gear in those racks ². Some stages behind that \$6,400 are expensive because they are hard to make, some because capacity is short, and some because one firm sets the price.

How it works

Three numbers set the cost of a die, the finished rectangle of circuitry cut out of a wafer.

- **Wafer price.** TSMC does not publish it. TSMC makes Nvidia's Blackwell chips on a custom-built process called 4NP ⁴. Epoch AI estimates \$17,000 for a 12-inch (300 mm) wafer on that process ¹.
- **Dies per wafer.** A wafer is a circle and dies are rectangles, so the edge of the wafer goes to waste. The bigger the die, the fewer fit. Each of the two logic dies

¹ NVIDIA's B200 costs around \$6,400 to produce. Epoch AI · 10 December 2025 [\[A\]](https://epoch.ai/data-insights/b200-cost-breakdown) epoch.ai/data-insights/b200-cost-breakdown

⁵ SK hynix Announces 2Q26 Financial Results. SK hynix · 29 July 2026 [\[A\]](https://news.skhynix.com/en/q2-2026-business-results/) news.skhynix.com/en/q2-2026-business-results/

¹¹ Micron Technology, Inc. Reports Record Fiscal Fourth-Quarter and Full-Year 2026 Results. Micron Technology · 30 September 2026 [\[A\]](https://investors.micron.com/news/press-release/2026/Micron-Technology-Inc--Reports-Record-Fiscal-Fourth-Quarter-and-Full-Year-2026-Results/default.aspx) investors.micron.com/news/press-release/2026/Micron-Technology-Inc--Reports-Record-Fiscal-Fourth-Quarter-and-Full-Year-2026-Results/default.aspx

⁷ NVIDIA Announces Financial Results for Second Quarter Fiscal 2027. Nvidia · 26 August 2026 [\[A\]](https://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-second-quarter-fiscal-2027) nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-second-quarter-fiscal-2027

² AI chip sales documentation. Epoch AI · 8 January 2026 [\[A\]](https://epoch.ai/data/ai-chip-sales-documentation/methodology#nvidia) epoch.ai/data/ai-chip-sales-documentation/methodology#nvidia

⁴ NVIDIA Blackwell Architecture. Nvidia · 18 December 2025 [\[A\]](https://www.nvidia.com/en-us/data-center/technologies/blackwell-architecture/) www.nvidia.com/en-us/data-center/technologies/blackwell-architecture/

in a B200 is about 800 mm²¹. Nvidia calls them reticle-limited, meaning they are near the largest size a scanner can print in one shot ⁴.

- **Yield.** Yield is the share of dies that work. Defects land at random, so a bigger die is more likely to catch one. Epoch AI models Blackwell die yield at 40 to 70 percent, centered on 60 ¹.

Together, those three numbers put the two dies inside a B200 at about \$900 ¹. The estimate counts only variable costs and leaves out fixed costs such as research and development ¹.

Variants and trade-offs

Memory is the largest line

The largest line on the bill is the 192 GB of high-bandwidth memory beside those dies, a generation called HBM3E ¹. At \$14 to \$17 per gigabyte it comes to about \$2,900, which Epoch AI counts as half of the build cost ¹. Epoch AI says HBM prices are negotiated deal by deal and can vary significantly with volume and contract length ¹. The \$14 to \$17 range dates from December 2025, before the 2026 rise in memory makers' sales and margins described below ⁵.

Three firms, SK hynix, Samsung and Micron, dominate the world market for that memory ³. SK hynix said in July 2026 that demand for its products had exceeded its available supply in recent quarters ⁶.

Packaging, substrate and test

CoWoS-L, the TSMC process that mounts dies and memory on one carrier and wires them together, adds about \$1,100, more than the logic dies it carries ¹. Epoch AI models packaging yield at 65 to 95 percent ¹. A package that fails testing loses its logic dies and memory stacks ¹. Spreading that loss over the good units adds about \$1,000 to each B200 ¹. Power delivery hardware, the circuit board, and final assembly and testing add \$480 ¹. See Advanced packaging for why packaging capacity runs out first.

The B200 bill of materials

LINE	ESTIMATE	RANGE
------	----------	-------

1 NVIDIA's B200 costs around \$6,400 to produce. Epoch AI · 10 December 2025 [\[A\]](#) epoch.ai/data-insights/b200-cost-breakdown

4 NVIDIA Blackwell Architecture. Nvidia · 18 December 2025 [\[A\]](#) www.nvidia.com/en-us/data-center/technologies/blackwell-architecture/

5 SK hynix Announces 2Q26 Financial Results. SK hynix · 29 July 2026 [\[A\]](#) news.skhynix.com/en/q2-2026-business-results/

3 Understanding the Biden Administration's Updated Export Controls. Center for Strategic and International Studies · 11 December 2024 [\[A\]](#) www.csis.org/analysis/understanding-biden-administrations-updated-export-controls


6 SK hynix Inc., Form 424B4 prospectus for the period ended 2026-07-10 (424(B)(4)). U.S. Securities and Exchange Commission (filing by SK hynix Inc.) · 10 July 2026 [\[A\]](#) www.sec.gov/Archives/edgar/data/2120882/000119312526299963/d32785d424b4.htm

HBM3E, 192 GB	\$2,900	\$2,800-3,100
CoWoS-L packaging	\$1,100	\$1,000-1,200
Packaging yield loss	\$1,000	\$430-1,700
Logic die, two on TSMC 4NP	\$900	\$720-1,200
Power delivery, board, assembly and test	\$480	\$370-600
Manufacturing cost	~\$6,400	\$5,700-7,300
Reported selling price per chip, March 2024	\$30,000-40,000	
Nvidia compute revenue per GPU in GB200 racks, data to October 2025	\$33,000-42,000	

Epoch AI models each line as a range, deriving the packaging line from TSMC's advanced packaging revenue and Nvidia's share of CoWoS capacity ¹. The estimate covers the accelerator module only and leaves out servers and rack-scale systems ¹. Epoch AI notes that most Blackwell revenue comes from those systems, which may carry lower margins ¹. So, Epoch AI says, Nvidia's actual margin on Blackwell may be lower than the gap between this cost and the reported chip price suggests ¹.

The B200 is used here because Epoch AI's line-by-line estimate covers it. In its results for the quarter to July 2026, Nvidia said its Vera Rubin platform was ramping into full production ⁷. Each Rubin GPU carries up to 288 GB of HBM4, a newer generation of the stacked memory ⁸.

Where the ~\$6,400 of B200 manufacturing cost goes, Epoch AI estimate, December 2025 \$



Source: Epoch AI B200 cost model, from a \$17,000 4NP wafer, 40–70% die yield, \$14–17 per GB of memory, and a packaging line derived from TSMC's advanced packaging revenue

From package to rack to cluster

At rack scale, cost is measured per GPU-hour, one GPU running for one hour. For an operator that owns the machine, SemiAnalysis put a Vera Rubin NVL72 at \$3.57

1 NVIDIA's B200 costs around \$6,400 to produce. Epoch AI · 10 December 2025 epoch.ai/data-insights/b200-cost-breakdown

7 NVIDIA Announces Financial Results for Second Quarter Fiscal 2027. Nvidia · 26 August 2026 nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-second-quarter-fiscal-2027

8 Inside NVIDIA Rubin GPU Architecture: Powering the Era of Agentic AI. Nvidia · 21 July 2026 developer.nvidia.com/blog/inside-nvidia-rubin-gpu-architecture-powering-the-era-of-agentic-ai/

per GPU-hour in July 2026, against \$1.84 for a GB200 and \$2.36 for a GB300 ⁹. Each generation costs its owner more per hour than the last ⁹.

But SemiAnalysis found Rubin cheaper per token than GB200 and GB300 at every interactivity level it compared ⁹. A token is a short chunk of text, and interactivity is how many tokens per second each user receives. That finding rests on early results from the cloud provider CoreWeave running the DeepSeek R1 model, and SemiAnalysis said it had not verified them ⁹.

One level up, SemiAnalysis estimated in June 2024 that a cluster of 100,000 of Nvidia's earlier H100 GPUs cost over \$4 billion in servers ¹⁰. It needs more than 150 MW of data center capacity and uses 1.59 TWh a year, which costs \$123.9 million at \$0.078 per kWh ¹⁰. In the network design Nvidia recommended, 98,304 optical transceivers link its GPUs to the first layer of switches, called leaf switches ¹⁰.

Who captures the value

FIRM	STAGE	QUARTER ENDED	Gross margin
Micron	HBM and DRAM	September 2026	86.8% ¹¹
Nvidia	Accelerator design	July 2026	75.0% ⁷
TSMC	Wafers and CoWoS	September 2026	66.4% ¹²
KLA (non-GAAP)	Inspection	June 2026	62.4% ¹³
ASML	Lithography	September 2026	57.5% ¹⁴
Lam Research	Etch	June 2026	51.7% ¹⁵

⁹ Vera Rubin NVL72 vs GB200 NVL72? Inference TCO & Architecture Analysis. SemiAnalysis · 23 July 2026 ^[B]
inference.semianalysis.com/blog/vera-rubin-nvl72-vs-gb200-nvl72-inference

¹⁰ 100,000 H100 Clusters: Power, Network Topology, Ethernet vs InfiniBand, Reliability, Failures, Checkpointing. SemiAnalysis · 17 June 2024 ^[B] newsletter.semianalysis.com/p/100000-h100-clusters-power-network

¹¹ Micron Technology, Inc. Reports Record Fiscal Fourth-Quarter and Full-Year 2026 Results. Micron Technology · 30 September 2026 ^[A] investors.micron.com/news/press-release/2026/Micron-Technology-Inc--Reports-Record-Fiscal-Fourth-Quarter-and-Full-Year-2026-Results/default.aspx

⁷ NVIDIA Announces Financial Results for Second Quarter Fiscal 2027. Nvidia · 26 August 2026 ^[A]
nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-second-quarter-fiscal-2027

¹² TSMC Reports Third Quarter EPS of NT\$29.43. U.S. Securities and Exchange Commission (filing by TAIWAN SEMICONDUCTOR MANUFACTURING CO LTD) · 15 October 2026 ^[A]
www.sec.gov/Archives/edgar/data/1046179/000104617926000726/a3q26e_withguidancexfinal.htm

¹³ Letter to Shareholders Q4 Fiscal 2026. KLA Corporation · 28 July 2026 · p. 3 ^[A]
dl1io3yog0oux5.cloudfront.net/_b9ee755a5f60dd0fb3f9e27967aed6af/klatencor/db/1117/10668/letter_to_shareholders/KLA+Earnings+Shareholder+Letter+-+Q4+FY26.pdf#page=3

¹⁴ ASML reports €11.7 billion total net sales and €5.0 billion net income in Q3 2026. U.S. Securities and Exchange Commission (filing by ASML HOLDING NV) · 14 October 2026 ^[A]
www.sec.gov/Archives/edgar/data/937966/000162828026065813/pressreleasefinancialresul.htm

¹⁵ Lam Research Corporation Reports Financial Results for the Quarter Ended June 28, 2026. Lam Research · 29 July 2026 ^[A]
newsroom.lamresearch.com/2026-07-29-Lam-Research-Corporation-Reports-Financial-Results-for-the-Quarter-Ended-June-28,-2026

Applied Materials	Deposition	July 2026	50.3% ¹⁶
ASE Technology	Packaging and test segment	June 2026	27.3% ¹⁷
Amkor	OSAT	June 2026	16.8% ¹⁸

Each figure except KLA's comes from the firm's own earnings release. KLA's shareholder letter gives only a non-GAAP gross margin for the quarter, 62.4 percent ¹³. KLA says its non-GAAP figures exclude certain costs and expenses that it believes do not reflect its core operating results ¹⁹. For ASE the table shows its packaging and test segment ¹⁷. The company as a whole had a gross margin of 21.0 percent, because about a third of its revenue came from electronics manufacturing services, which earned 8.9 percent ¹⁷.

SK hynix's earnings release gives no gross margin. Its operating margin was 76 percent in the quarter to June 2026, and its revenue was up 257 percent from a year earlier ⁵. The substrate maker Ibiden gives no gross margin either, but its operating profit was 21.8 percent of net sales in the quarter to June 2026 ²⁰.

16 Applied Materials Announces Third Quarter 2026 Results. Applied Materials · 13 August 2026 [A](https://ir.appliedmaterials.com/news-releases/news-release-details/applied-materials-announces-third-quarter-2026-results)

17 ASE Technology Holding Co., Ltd. Reports Its Unaudited Consolidated Financial Results for the Second Quarter of 2026. PR Newswire · 30 July 2026 [A](https://www.prnewswire.com/news-releases/ase-technology-holding-co-ltd-reports-its-unaudited-consolidated-financial-results-for-the-second-quarter-of-2026-302838714.html)


18 AMKOR TECHNOLOGY, INC., Form 8-K current report for the period ended 2026-07-27 (8-K). U.S. Securities and Exchange Commission (filing by AMKOR TECHNOLOGY, INC.) · 27 July 2026 [A](https://www.sec.gov/Archives/edgar/data/0001047127/000104712726000043/amkr6302026erex-991.htm)

13 Letter to Shareholders Q4 Fiscal 2026. KLA Corporation · 28 July 2026 · p. 3 [A](https://d1io3yogOoux5.cloudfront.net/_b9ee755a5f60dd0fb3f9e27967aed6af/klatencor/db/1117/10668/letter_to_shareholders/KLA+Earnings+Shareholder+Letter+-+Q4+FY26.pdf#page=3)

19 KLA Corporation Reports Fiscal 2026 Fourth Quarter and Full Year Results. KLA Corporation · 28 July 2026 [A](https://ir.kla.com/news-events/press-releases/detail/518/kla-corporation-reports-fiscal-2026-fourth-quarter-and-full)

5 SK hynix Announces 2Q26 Financial Results. SK hynix · 29 July 2026 [A](https://news.skhyunix.com/en/q2-2026-business-results/)

20 Outline of Consolidated Financial Results For the First Quarter Ended June 30, 2026. Ibiden · 4 August 2026 [A](https://www.ibiden.com/ir/items/outline2026Q1.pdf)



Source: Each firm's own quarterly earnings release; quarters end between June and September 2026; ASE's bar is its packaging and test segment; KLA's bar is the non-GAAP figure from its shareholder letter (TSMC shown)

In the table, the firms in design, equipment and leading-edge wafers earn gross margins of 50.3 to 75.0 percent, and those in assembly and test 16.8 to 27.3 percent. In 2026 Micron passed all of them. Its gross margin rose from 39.8 percent for the year to August 2025 to 86.8 percent in the quarter to September 2026 ¹¹.

The Semiconductor Industry Association, using monthly data compiled by World Semiconductor Trade Statistics, counted \$791.7 billion of global chip sales in 2025 ²¹. In its June 2026 forecast, World Semiconductor Trade Statistics put 2025 at \$795.6 billion and 2026 at \$1.51 trillion, with memory growing about 250 percent and logic 37 percent ²².

The industry group SEMI counted \$135.1 billion of semiconductor equipment billings in 2025 ²³. Capital spending at Alphabet, Amazon, Meta, Microsoft and Oracle came to \$448.2 billion in 2025, the sum of Epoch AI's quarterly figures ²⁴. Epoch AI says that spending would reach \$770 billion in 2026 if the trend continues ²⁴. These totals overlap, so they cannot be added together.

11 Micron Technology, Inc. Reports Record Fiscal Fourth-Quarter and Full-Year 2026 Results. Micron Technology · 30 September 2026 [\[A\]](https://investors.micron.com/news/press-release/2026/Micron-Technology-Inc--Reports-Record-Fiscal-Fourth-Quarter-and-Full-Year-2026-Results/default.aspx) investors.micron.com/news/press-release/2026/Micron-Technology-Inc--Reports-Record-Fiscal-Fourth-Quarter-and-Full-Year-2026-Results/default.aspx


21 Global Annual Semiconductor Sales Increase 25.6% to \$791.7 Billion in 2025. Semiconductor Industry Association · 6 February 2026 [\[A\]](https://www.semiconductors.org/global-annual-semiconductor-sales-increase-25-6-to-791-7-billion-in-2025/) www.semiconductors.org/global-annual-semiconductor-sales-increase-25-6-to-791-7-billion-in-2025/

22 Global Semiconductor Market Surges Beyond USD 1.5 Trillion in 2026 Driven by Extraordinary Memory Expansion. World Semiconductor Trade Statistics · 2 June 2026 [\[A\]](https://www.wsts.org/esraCMS/extension/media/f/WST/7618/WSTS_FC-Release-2026-May.pdf) www.wsts.org/esraCMS/extension/media/f/WST/7618/WSTS_FC-Release-2026-May.pdf

23 SEMI Reports Global Semiconductor Equipment Billings Reached \$135 Billion in 2025, Up 15% Year-on-Year. SEMI · 7 April 2026 [\[A\]](https://www.semi.org/en/SEMI-Reports-Global-Semiconductor-Equipment-Billings-Reached-135-Billion-in-2025) www.semi.org/en/SEMI-Reports-Global-Semiconductor-Equipment-Billings-Reached-135-Billion-in-2025

24 Hyperscaler capex has quadrupled since GPT-4's release. Epoch AI · 26 February 2026 [\[A\]](https://epoch.ai/data-insights/hyperscaler-capex-trend) epoch.ai/data-insights/hyperscaler-capex-trend

Selected market totals, 2025 \$B



Source: Semiconductor Industry Association on chip sales; Epoch AI quarterly capital spending at Alphabet, Amazon, Meta, Microsoft and Oracle, summed for 2025; Nvidia data center revenue for its fiscal year to January 2026; SEMI on equipment billings

Where the money goes

In Epoch AI's estimate for the B200, packaging cost \$1,100 and packaging yield loss another \$1,000, against \$900 for the logic dies ¹.

The table below sorts each stage by whether its cost falls as volume grows. The classes are a judgment drawn from the sources in this chapter and the stage chapters.

STAGE	CLASS	WHY
EUV lithography	Structurally expensive	One supplier, ASML. See Lithography.
Masks and pellicles	Structurally expensive	Every design needs its own mask set, and every EUV layer a flawless mask. See Photomasks and pellicles.
Metrology and inspection	Structurally expensive	KLA held 56.8 percent of the market for these tools in 2025, by CSET's analysis of TechInsights data ²⁵ . See Metrology and inspection.
HBM	Structurally expensive	SK hynix says HBM needs more wafer input than ordinary DRAM ⁶ . See Memory and HBM.
Advanced packaging, CoWoS-L	Structurally expensive	Modeled packaging yield is 65 to 95 percent, and a failed package loses its dies and memory.
ABF substrates	Structurally expensive	Investment can add substrate lines, but Ajinomoto said in June 2026 that the insulating film it makes for these substrates has held over 95 percent of its market since launch ²⁶ . See Substrates and PCBs.
Power and cooling plant	Structurally expensive	In the United States, grid connections face long queues ²⁷ . GE Vernova's gas equipment backlog and slot reservations reached 116 GW in mid-2026 ²⁸ . See Data centers and power.

1 NVIDIA's B200 costs around \$6,400 to produce. Epoch AI · 10 December 2025 [\[A\]](#) epoch.ai/data-insights/b200-cost-breakdown

25 Supply Chain Explorer: Advanced Chips. Emerging Technology Observatory (ETO) · Accessed 9 September 2026 [\[A\]](#) chipexplorer.eto.tech/?parentNode=N118&selectedNode=N60

6 SK hynix Inc., Form 424B4 prospectus for the period ended 2026-07-10 (424(B)(4)). U.S. Securities and Exchange Commission (filing by SK hynix Inc.) · 10 July 2026 [\[A\]](#) www.sec.gov/Archives/edgar/data/2120882/000119312526299963/d32785d424b4.htm

26 Ajinomoto, Growth Strategy for the Electronic Materials Business, 30 June 2026. Ajinomoto · 30 June 2026 · p. 3 [\[A\]](#) www.ajinomoto.co.jp/company/en/ir/event/business_briefing/main/011117/teaserItems1/01/linkList/03/link/20260630_presentation_E.pdf#page=3

27 OpenAI Stargate: where the US sites stand. Epoch AI · 17 April 2026 [\[A\]](#) epoch.ai/publications/openai-stargate-where-the-us-sites-stand

28 GE Vernova Inc., Form 8-K current report for the period ended 2026-07-22 (8-K). U.S. Securities and Exchange Commission (filing by GE Vernova Inc.) · 22 July 2026 [\[A\]](#) www.sec.gov/Archives/edgar/data/1996810/000199681026000147/gevpressrelease2q26.htm

Leading-edge logic wafers	Compressible with scale	A very large die yields well only once a process matures. See Transistors and the front end.
Design and EDA	Compressible with scale	A design is paid for once, and its cost is spread across every chip sold. See Chip design, EDA and IP.
Deposition and etch tools	Compressible with scale	Most deposition and etch steps have three or four established tool makers. See Deposition and etch.
Optical transceivers	Compressible with scale	Nvidia agreed in March 2026 to invest \$2 billion in the laser maker Lumentum, with rights to future capacity ²⁹ . It agreed a similar \$2 billion deal with Coherent ³⁰ . See Systems and networking.
Test and assembly (OSAT)	Already commoditized	Gross margins of 16.8 to 27.3 percent at Amkor and ASE's packaging and test segment, in the table above.
Wafers and bulk gases	Already commoditized	Several firms make approved 300 mm wafers. By Siltronic's estimate, the five largest wafer makers served around 75 percent of the market in the first three quarters of 2025 ³¹ . See Silicon and wafers.
Rack integration (ODM)	Already commoditized	In October 2025 Nvidia said more than 50 makers of systems and components were gearing up for the Vera Rubin NVL72 ³² . See Systems and networking.

THREE THINGS TO REMEMBER

- Epoch AI estimated in December 2025 that a B200 cost about \$6,400 to build, against a selling price of \$30,000 to \$40,000 per chip reported in March 2024.
- The stacked memory is the most expensive part, about half of the cost.
- Packaging costs more than the logic chips, and each package that fails wastes the chips and memory inside it.

Made by **Yashvardhan Sharma** with Claude. Corrections and questions: contact@chipsupplychain.org.

²⁹ NVIDIA Announces Strategic Partnership With Lumentum to Develop State-of-the-Art Optics Technology. Nvidia · 2 March 2026 [\[A\]](https://nvidianews.nvidia.com/news/nvidia-announces-strategic-partnership-with-lumentum-to-develop-state-of-the-art-optics-technology) nvidianews.nvidia.com/news/nvidia-announces-strategic-partnership-with-lumentum-to-develop-state-of-the-art-optics-technology

³⁰ NVIDIA and Coherent Announce Strategic Partnership to Develop Optics Technology to Scale Next-Generation Data Center Architecture. Nvidia · 2 March 2026 [\[A\]](https://nvidianews.nvidia.com/news/nvidia-and-coherent-announce-strategic-partnership-to-develop-optics-technology-to-scale-next-generation-data-center-architecture) nvidianews.nvidia.com/news/nvidia-and-coherent-announce-strategic-partnership-to-develop-optics-technology-to-scale-next-generation-data-center-architecture

³¹ Investor Presentation: Foundation of Digital Life. Siltronic · April 2026 · p. 6 [\[A\]](https://www.siltronic.com/fileadmin/investorrelations/2026/Q1/20260429_Siltronic_InvestorPresentation_.pdf#page=6) www.siltronic.com/fileadmin/investorrelations/2026/Q1/20260429_Siltronic_InvestorPresentation_.pdf#page=6

³² NVIDIA, Partners Drive Next-Gen Efficient Gigawatt AI Factories in Buildup for Vera Rubin. Nvidia · 13 October 2025 [\[A\]](https://blogs.nvidia.com/blog/gigawatt-ai-factories-ocp-vera-rubin) blogs.nvidia.com/blog/gigawatt-ai-factories-ocp-vera-rubin

UPDATED 20 OCTOBER 2026

A company filings, governments, standards bodies, universities and research institutions; **B** SemiAnalysis, Reuters, the Financial Times and Nikkei Asia.

HOW TO CITE THIS CHAPTER

Sharma, Yashvardhan. "The Economics of AI Chips." Chip Supply Chain: How AI Chips Are Made. Edition of 20 October 2026. <https://chipsupplychain.org/component-costs/>

```
@misc{chipsupplychain2026componentcosts,  
  title      = {The Economics of AI Chips},  
  author     = {Sharma, Yashvardhan},  
  howpublished = {Chip Supply Chain: How AI Chips Are Made},  
  year       = {2026},  
  month      = {October},  
  note       = {Edition of 20 October 2026},  
  url        = {https://chipsupplychain.org/component-costs/}  
}
```