

Chapter 14: Systems and Networking

chipsupplychain.org/systems-and-networking/

CITE AS Sharma, Yashvardhan. "Systems and Networking." *Chip Supply Chain: How AI Chips Are Made*. Edition of 20 October 2026. <https://chipsupplychain.org/systems-and-networking/>

CHAPTER

14

STAGE 14 OF 15

Systems and Networking

AI chips are wired together 72 to a rack and run as one computer. Many firms can assemble the racks, but lasers for the optical links between them were scarce in 2026.

IN PLAIN TERMS

One AI chip is not enough to train a model. So seventy-two of them are wired together in a rack, a cabinet the size of a wardrobe, and run as one computer. Each chip sits on its own board and draws more than a thousand amps, several times the current a whole house is wired to carry. The chips have to swap results with each other all the time. A copper wire can carry those signals only a few meters before they fade, so the longer links send them as light down glass fiber. Many companies can build the rack, but in 2026 the lasers that send the light were scarce.

IN SHORT

In March 2026 Nvidia agreed to invest \$2 billion in the laser maker Lumentum, alongside a multibillion-dollar purchase commitment and rights to future capacity for advanced laser components ³. Indium phosphide is used in high-speed lasers for optical communications, and China produced 70 percent of the world's indium in 2025 ¹³.

On 4 February 2025 China added indium phosphide substrates, the wafers the lasers are grown on, to its export control list ¹⁴. AXT, a US-listed maker of those wafers, does all its manufacturing in China ¹⁴. In its filing for the quarter to June 2026, AXT said it could not estimate when it would resume shipping the wafers to the United States ¹⁴.

³ 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

¹³ Mineral Commodity Summaries 2026: Indium. U.S. Geological Survey · 5 February 2026 [\[A\]](https://pubs.usgs.gov/periodicals/mcs2026/mcs2026-indium.pdf) pubs.usgs.gov/periodicals/mcs2026/mcs2026-indium.pdf

¹⁴ AXT INC, Form 10-Q quarterly report for the period ended 2026-06-30. U.S. Securities and Exchange Commission (filing by AXT INC) · 13 August 2026 [\[A\]](https://www.sec.gov/Archives/edgar/data/1051627/000143774926027677/axti20260630_10q.htm) www.sec.gov/Archives/edgar/data/1051627/000143774926027677/axti20260630_10q.htm

CHOKEPOINT CARD		HIGH CONCENTRATION
		MODERATELY HARD TO SUBSTITUTE
CONCENTRATION	SUBSTITUTABILITY	
High	Moderate	
	Rated on NVLink and its NVSwitch chips, which Nvidia controls and which join the accelerators in a rack ¹ . Lasers for the optical links were also scarce in 2026 ² . Nvidia secured future capacity rights for them from Lumentum in March 2026 ³ .	
	PRICE OR MARKET SIZE	
	\$3.1 million	
	for one GB200 NVL72 rack bought by a typical hyperscaler, or about \$3.9 million with networking, storage and other items, by SemiAnalysis's August 2025 estimate ⁴	
WHO LEADS	US Nvidia NVLink, InfiniBand and Spectrum-X Ethernet, which it credited for data center growth in the quarter to April 2026⁵ US Broadcom \$16.7B of AI semiconductor revenue in the quarter to August 2026, up 221%⁶ TW Foxconn World's largest AI server provider, its chairman said in March 2026⁷ CN InnoLight High-speed optical modules made in Suzhou and, since 2019, Thailand⁸ US Coherent Lasers and optical networking products, with a \$2B Nvidia investment agreed in March 2026⁹	
WHERE IT IS MADE	US United States Switch and accelerator silicon, system design, indium phosphide lasers TW Taiwan Contract design, rack assembly and high-layer-count circuit boards JP Japan Low-loss board laminate CN China Optical module assembly and test, and most of the world's indium TH Thailand Chinese module makers' offshore transceiver plants	
WHY SUBSTITUTION IS POSSIBLE	UALink is an open specification for the same job as NVLink ¹⁰ . AMD's Helios rack, launched in July 2026, runs it over Ethernet ¹¹ . Since May 2025 Nvidia's NVLink Fusion program has let other firms' custom chips use NVLink, so the dependence on Nvidia stays ¹² .	
WHERE CHINA STANDS	WHERE THE US STANDS	
Chinese firms such as InnoLight assemble high-speed optical modules in volume, at sites including Suzhou and, since 2019, Thailand ⁸ . In February 2025 China put indium under new export restrictions ¹³ . The same month it began requiring export permits for indium phosphide substrates, the wafers the lasers are grown on ¹⁴ .	Coherent broke ground in Sherman, Texas, in June 2026 on an expansion of its indium phosphide factory ¹⁵ . In August 2026 Nvidia said that Wistron's new plant in Fort Worth, Texas, was producing its GB300 superchips and that Foxconn was building a factory in Houston ¹⁶ .	

1 NVIDIA CORP, Form 10-K annual report for the period ended 2026-01-25 (10-K) U.S. Securities and Exchange Commission (filing by NVIDIA CORP) · 25 February 2026 www.sec.gov/Archives/edgar/data/1045810/000104581026000021/nvda-20260125.htm

Before a chip package can compute, it has to be built into a machine. Building that machine is a separate industry from making the package.

How it works

An accelerator (the AI chip) ships as a module, the chip package on a small board with the circuits that feed it power. Nvidia calls its version SXM, and the open-standards equivalent is the Open Compute Project's accelerator module. Eight modules bolt onto a baseboard that carries their power and the links between them. On Nvidia's baseboard, switch chips pass data between any two modules at full speed.

Rack-scale systems drop the baseboard and put the chips on trays stacked in the rack. An Nvidia GB300 NVL72 holds 18 compute trays plus nine switch trays, 72 GPUs in all, and draws up to 142 kW¹⁷. The newer Vera Rubin NVL72 uses the same rack footprint¹⁸. SemiAnalysis estimated in August 2025 that a GB200 NVL72

-
- 2 AI demand strains supplies of lasers, fiber and other optical tech. Nikkei Asia · 27 May 2026 [\[B\]](#)
asia.nikkei.com/business/technology/tech-asia/ai-demand-strains-supplies-of-lasers-fiber-and-other-optical-tech
 - 3 NVIDIA Announces Strategic Partnership With Lumentum to Develop State-of-the-Art Optics Technology. Nvidia · 2 March 2026 [\[A\]](#) [nvidianews.nvidia.com/news/nvidia-announces-strategic-partnership-with-lumentum-to-develop-state-of-the-art-optics-technology](#)
 - 4 H100 vs GB200 NVL72 Training Benchmarks - Power, TCO, and Reliability Analysis, Software Improvement Over Time. SemiAnalysis · 20 August 2025 [\[B\]](#) [newsletter.semianalysis.com/p/h100-vs-gb200-nvl72-training-benchmarks](#)
 - 5 NVIDIA CORP, Form 10-Q quarterly report for the period ended 2026-04-26 (10-Q). U.S. Securities and Exchange Commission (filing by NVIDIA CORP) · 20 May 2026 [\[A\]](#) [www.sec.gov/Archives/edgar/data/1045810/000104581026000052/nvda-20260426.htm](#)
 - 6 Broadcom Inc., Form 8-K current report for the period ended 2026-09-02 (8-K). U.S. Securities and Exchange Commission (filing by Broadcom Inc.) · 2 September 2026 [\[A\]](#) [www.sec.gov/Archives/edgar/data/1730168/000173016826000076/avgo-08022026x8kxex99.htm](#)
 - 7 Hon Hai Technology Group (Foxconn) Announces FY2025 & 4Q25 Financial Results. Hon Hai Precision Industry (Foxconn) · 16 March 2026 [\[A\]](#) [www.honhai.com/en-us/press-center/press-releases/latest-news/1978](#)
 - 8 InnoLight Technology, company overview page. InnoLight Technology · Accessed 9 September 2026 [\[A\]](#) [www.innolight.com/about](#)
 - 9 NVIDIA and Coherent Announce Strategic Partnership to Develop Optics Technology to Scale Next-Generation Data Center Architecture. Nvidia · 2 March 2026 [\[A\]](#) [nvidianews.nvidia.com/news/nvidia-and-coherent-announce-strategic-partnership-to-develop-optics-technology-to-scale-next-generation-data-center-architecture](#)
 - 10 UALink™ 200G 1.0 Specification Overview. UALink Consortium · 15 May 2025 [\[A\]](#) [ualinkconsortium.org/blog/ualink-200g-1-0-specification-overview-802/](#)
 - 11 AAI 2026: AMD Launches AMD Helios Rackscale Solution for Frontier AI. AMD Newsroom · 23 July 2026 [\[A\]](#)
newsroom.amd.com/news/aai-2026-helios-update/
 - 12 NVIDIA Unveils NVLink Fusion for Industry to Build Semi-Custom AI Infrastructure With NVIDIA Partner Ecosystem. Nvidia · 18 May 2025 [\[A\]](#) [nvidianews.nvidia.com/news/nvidia-nvlink-fusion-semi-custom-ai-infrastructure-partner-ecosystem](#)
 - 13 Mineral Commodity Summaries 2026: Indium. U.S. Geological Survey · 5 February 2026 [\[A\]](#)
pubs.usgs.gov/periodicals/mcs2026/mcs2026-indium.pdf
 - 14 AXT INC, Form 10-Q quarterly report for the period ended 2026-06-30. U.S. Securities and Exchange Commission (filing by AXT INC) · 13 August 2026 [\[A\]](#) [www.sec.gov/Archives/edgar/data/1051627/000143774926027677/axti20260630_10q.htm](#)
 - 15 Coherent Announces a CHIPS Letter of Intent for \$50 Million to Expand World-Leading Manufacturing Facility for AI Infrastructure. Coherent · 16 June 2026 [\[A\]](#) [www.coherent.com/news/press-releases/a-chip-letter-of-intent-for-50m-to-expand-world-leading-manufacturing-facility-for-ai-infrastructure](#)
 - 16 NVIDIA and Partners Build in America, for America. Nvidia · 5 August 2026 [\[A\]](#) [blogs.nvidia.com/blog/nvidia-and-partners-build-in-america-for-america/](#)
 - 17 System Hardware & Components. Nvidia · Accessed 9 September 2026 [\[A\]](#) [docs.nvidia.com/enterprise-reference-architectures/nvl72-ai-factory/latest/components.html](#)
 - 18 NVIDIA, Partners Drive Next-Gen Efficient Gigawatt AI Factories in Buildup for Vera Rubin. Nvidia · 13 October 2025 [\[A\]](#)
blogs.nvidia.com/blog/gigawatt-ai-factories-ocp-vera-rubin

rack cost a typical hyperscaler \$3.1 million, or about \$3.9 million with networking, storage and other items ⁴.

Inside the rack the trays reach each other over copper ¹⁷. Between neighboring racks, the signal can travel up to 7 meters over active electrical cables, a type of copper cable sold by firms such as Credo ¹⁹. Longer links run on glass fiber, with an optical transceiver at each end. A laser in the sending transceiver turns the electrical signal into pulses of light, and a detector in the receiving one turns the pulses back into current.

The accelerators in an AI cluster exchange data over two networks that are not interchangeable. Scale-up connects anywhere from a few dozen accelerators to thousands closely enough that they can read each other's memory. It runs at terabytes per second. Scale-out joins those groups to one another at hundreds of gigabits per second. A byte is eight bits, so each GPU in a GB300 NVL72 has nine times as much bandwidth for scale-up as for scale-out, SemiAnalysis calculated in January 2026 ²⁰.

A model too big for one accelerator is split across many. The parts that exchange data constantly stay inside the scale-up group, while whole blocks of layers can sit on different groups across the slower scale-out network ²¹. So scale-up bandwidth limits how efficiently a large model trains.

Variants and trade-offs

Scale-up: NVLink, UALink and Google's torus

The designs differ in how many accelerators can share one pool of memory.

- **NVLink and NVSwitch.** Nvidia's per-GPU bandwidth went from 900 GB/s in NVLink 4 to 1,800 in NVLink 5, and the rack's own network moves 130 TB/s in Blackwell ²². Nvidia's Rubin launch release of January 2026 gives NVLink 6 as 3,600 GB/s per GPU and 260 TB/s for a Vera Rubin NVL72 rack ²³. As of October 2026, Nvidia's NVLink product page still lists preliminary figures of 3,000 GB/s and 216 TB/s ²².
- **UALink.** The open specification runs 200 Gb/s per lane, four lanes to a station for 800 Gb/s, and can address up to 1,024 endpoints in one group ¹⁰. AMD's

4 H100 vs GB200 NVL72 Training Benchmarks - Power, TCO, and Reliability Analysis, Software Improvement Over Time. SemiAnalysis · 20 August 2025 [B] newsletter.semianalysis.com/p/h100-vs-gb200-nvl72-training-benchmarks

17 System Hardware & Components. Nvidia · Accessed 9 September 2026 [A] docs.nvidia.com/enterprise-reference-architectures/nvl72-ai-factory/latest/components.html

19 Credo Technology Group Holding Ltd, Form 10-K annual report for the period ended 2026-05-02 (10-K). U.S. Securities and Exchange Commission (filing by Credo Technology Group Holding Ltd) · 15 June 2026 [A] www.sec.gov/Archives/edgar/data/1807794/000162828026043303/crdo-20260502.htm

20 Co-Packaged Optics (CPO) Book – Scaling with Light for the Next Wave of Interconnect. SemiAnalysis · 1 January 2026 [B] newsletter.semianalysis.com/p/co-packaged-optics-cpo-book-scaling

21 Parallelisms Guide – Megatron Bridge. Nvidia · Accessed 9 September 2026 [A] docs.nvidia.com/nemo-framework/user-guide/latest/nemotoolkit/features/parallelisms.html

22 NVLink & NVLink Switch for Advanced Multi-GPU Communication. Nvidia · 20 April 2026 [A] www.nvidia.com/en-us/data-center/nvlink/


23 NVIDIA Kicks Off the Next Generation of AI With Rubin. Nvidia · 5 January 2026 [A] nvidianews.nvidia.com/news/rubin-platform-ai-supercomputer

10 UALink™ 200G 1.0 Specification Overview. UALink Consortium · 15 May 2025 [A] ualinkconsortium.org/blog/ualink-200g-1-0-specification-overview-802/

Helios rack, launched in July 2026, joins 72 GPUs and runs UALink over Ethernet for its scale-up network ¹¹.

- **Google's 3D torus.** Google wires 64 of its Ironwood TPU accelerators into each rack, called a cube ²⁴. Inside the cube each chip links directly to six others, in a three-dimensional grid that Google calls a 3D torus ²⁴. Optical circuit switches join cubes into superpods of 9,216 chips that can bypass a failed cube or link ²⁴.

NVLink bandwidth per GPU by generation GB/S




Source: NVIDIA NVLink product page (NVLink 4 and 5); NVIDIA Rubin launch release, January 2026 (NVLink 6)

Scale-out: InfiniBand versus Ethernet

InfiniBand, a specialized network long used in supercomputers, was the main scale-out network when AI clusters were small and uniform, but Ethernet switch capacity grew quickly. Broadcom's Tomahawk 4 switched 25.6 Tbps on one chip in December 2019 ²⁵. Tomahawk 5 reached 51.2 Tbps in August 2022, and Broadcom said it had doubled the bandwidth about every two years since 2014 ²⁶. Tomahawk 6 doubled it again to 102.4 Tbps in June 2025 ²⁷. In June 2025 the Ultra Ethernet Consortium published its 1.0 specification, a set of Ethernet-based communication protocols for AI and supercomputing that covers transport, congestion control and direct memory access ²⁸.

Broadcom Tomahawk switching capacity on one chip TBPS



Source: Broadcom product releases for Tomahawk 4, 5 and 6

11 AAI 2026: AMD Launches AMD Helios Rackscale Solution for Frontier AI. AMD Newsroom · 23 July 2026 newsroom.amd.com/news/aai-2026-helios-update/

24 Inside the Ironwood TPU codesigned AI stack. Google Cloud · 6 November 2025 cloud.google.com/blog/products/compute/inside-the-ironwood-tpu-codesigned-ai-stack

25 Broadcom Ships Tomahawk 4, Industry’s Highest Bandwidth Ethernet Switch Chip at 25.6 Terabits per Second. Broadcom · 9 December 2019 investors.broadcom.com/news-releases/news-release-details/broadcom-ships-tomahawk-4-industrys-highest-bandwidth-ethernet

26 Broadcom Ships Tomahawk 5, Industry's Highest Bandwidth Switch Chip to Accelerate AI/ML Workloads. Broadcom · 16 August 2022 investors.broadcom.com/news-releases/news-release-details/broadcom-ships-tomahawk-5-industrys-highest-bandwidth-switch

27 Broadcom Ships Tomahawk 6: World’s First 102.4 Tbps Switch. Broadcom · 3 June 2025 investors.broadcom.com/news-releases/news-release-details/broadcom-ships-tomahawk-6-worlds-first-1024-tbps-switch

28 Ultra Ethernet Consortium (UEC) Launches Specification 1.0 Transforming Ethernet for AI and HPC at Scale. Ultra Ethernet Consortium · 11 June 2025 ultraethernet.org/ultra-ethernet-consortium-uec-launches-specification-1-0-transforming-ethernet-for-ai-and-hpc-at-scale/

Nvidia sells switches and network cards for both InfiniBand and Ethernet. Its data center revenue reached \$89.0 billion in the quarter to July 2026, up 117 percent from a year earlier ²⁹. In the quarter before, it credited its data center growth to Blackwell 300 products and to demand for its "InfiniBand, Spectrum-X Ethernet, and NVLink" solutions ⁵.

Optics: pluggable, then co-packaged

A cluster needs several transceivers, the modules that turn electricity into light and back, for each accelerator. Their power cost is one reason Broadcom and Nvidia are moving the optics out of pluggable modules and into the switch chip's own package, a design called co-packaged optics. Broadcom's Tomahawk 6 Davisson cuts optical interconnect power by 70 percent compared with traditional pluggable modules ³⁰. It keeps its lasers in field-replaceable modules ³⁰. Nvidia's Spectrum-X Photonics switches, announced in March 2025, reach 512 ports of 800 gigabits and use four times fewer lasers than traditional designs ³¹. Nvidia builds them on its Spectrum-6 switch design ²³. In August 2026 Nvidia said its Spectrum-6 switch systems, with pluggable or co-packaged optics, were arriving at AI data centers around the world ³².

Who makes it

Chinese makers of optical modules have built plants outside China. InnoLight has run a plant in Thailand since 2019, alongside its sites in Suzhou and Taiwan ⁸. The

29 NVIDIA CORP, Form 10-Q quarterly report for the period ended 2026-07-26 (10-Q). U.S. Securities and Exchange Commission (filing by NVIDIA CORP) · 26 August 2026 [\[A\]](https://www.sec.gov/Archives/edgar/data/1045810/000104581026000075/nvda-20260726.htm) www.sec.gov/Archives/edgar/data/1045810/000104581026000075/nvda-20260726.htm

5 NVIDIA CORP, Form 10-Q quarterly report for the period ended 2026-04-26 (10-Q). U.S. Securities and Exchange Commission (filing by NVIDIA CORP) · 20 May 2026 [\[A\]](https://www.sec.gov/Archives/edgar/data/1045810/000104581026000052/nvda-20260426.htm) www.sec.gov/Archives/edgar/data/1045810/000104581026000052/nvda-20260426.htm

30 Broadcom Announces Tomahawk 6 – Davisson, the Industry's First 102.4-Tbps Ethernet Switch with Co-Packaged Optics. Broadcom · 8 October 2025 [\[A\]](https://investors.broadcom.com/news-releases/news-release-details/broadcom-announces-tomahawk-6-davisson-industrys-first-1024) investors.broadcom.com/news-releases/news-release-details/broadcom-announces-tomahawk-6-davisson-industrys-first-1024

31 NVIDIA Announces Spectrum-X Photonics, Co-Packaged Optics Networking Switches to Scale AI Factories to Millions of GPUs. Nvidia · 18 March 2025 [\[A\]](https://nvidianews.nvidia.com/news/nvidia-spectrum-x-co-packaged-optics-networking-switches-ai-factories) nvidianews.nvidia.com/news/nvidia-spectrum-x-co-packaged-optics-networking-switches-ai-factories

23 NVIDIA Kicks Off the Next Generation of AI With Rubin. Nvidia · 5 January 2026 [\[A\]](https://nvidianews.nvidia.com/news/rubin-platform-ai-supercomputer) nvidianews.nvidia.com/news/rubin-platform-ai-supercomputer

32 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

8 InnoLight Technology, company overview page. InnoLight Technology · Accessed 9 September 2026 [\[A\]](https://www.innolight.com/about) www.innolight.com/about

Thai subsidiary of Eoptolink, another Chinese module maker, lists sites in Chonburi and Rayong ³³.

The Federal Communications Commission publishes a Covered List of equipment and services judged to pose an unacceptable risk to US national security or to the security and safety of US persons ³⁴. Its July 2026 order bars new authorizations for devices that contain logic-bearing hardware components produced by firms on that list ³⁵. Commenters in the record named optical transceivers as one such component ³⁶. As of 16 October 2026, the Covered List did not include InnoLight or Coherent ³⁷.

AI data centers were straining supplies of the lasers inside those modules and pushing up their prices, Nikkei Asia reported in May 2026 ². Nvidia had agreed in March 2026 to invest \$2 billion in one laser maker, Lumentum, alongside a multibillion-dollar purchase commitment and future capacity access rights for advanced laser components ³. Coherent, another laser maker, broke ground in Sherman, Texas, in June 2026 on an expansion of its indium phosphide factory ¹⁵.

Indium phosphide is used in the high-speed laser diodes and photodetectors for optical communications ¹³. China accounted for 70 percent of world indium production in 2025 ¹³. In February 2025 China's Ministry of Commerce placed indium under new export restrictions ¹³.

33 Eoptolink Technology, contacts page. Eoptolink Technology · Accessed 9 September 2026 [A] www.eoptolink.com/about-us/contacts

34 Protecting Against National Security Threats to the Communications Supply Chain through the Equipment Authorization Program: Third Report and Order and Third Further Notice of Proposed Rulemaking (FCC-26-50). U.S. Federal Communications Commission · 23 July 2026 · p. 3 [A] docs.fcc.gov/public/attachments/FCC-26-50A1.pdf#page=3

35 Protecting Against National Security Threats to the Communications Supply Chain through the Equipment Authorization Program: Third Report and Order and Third Further Notice of Proposed Rulemaking (FCC-26-50). U.S. Federal Communications Commission · 23 July 2026 · p. 6 [A] docs.fcc.gov/public/attachments/FCC-26-50A1.pdf#page=6

36 Protecting Against National Security Threats to the Communications Supply Chain through the Equipment Authorization Program: Third Report and Order and Third Further Notice of Proposed Rulemaking (FCC-26-50). U.S. Federal Communications Commission · 23 July 2026 · p. 10 [A] docs.fcc.gov/public/attachments/FCC-26-50A1.pdf#page=10

37 List of Equipment and Services Covered By Section 2 of The Secure Networks Act. U.S. Federal Communications Commission · 16 October 2026 [A] www.fcc.gov/supplychain/coveredlist

2 AI demand strains supplies of lasers, fiber and other optical tech. Nikkei Asia · 27 May 2026 [B] asia.nikkei.com/business/technology/tech-asia/ai-demand-strains-supplies-of-lasers-fiber-and-other-optical-tech

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

15 Coherent Announces a CHIPS Letter of Intent for \$50 Million to Expand World-Leading Manufacturing Facility for AI Infrastructure. Coherent · 16 June 2026 [A] www.coherent.com/news/press-releases/a-chip-letter-of-intent-for-50m-to-expand-world-leading-manufacturing-facility-for-ai-infrastructure

13 Mineral Commodity Summaries 2026: Indium. U.S. Geological Survey · 5 February 2026 [A] pubs.usgs.gov/periodicals/mcs2026/mcs2026-indium.pdf

Broadcom's AI chips include custom accelerators and Ethernet switch chips ³⁸. In its annual report for fiscal 2025, Broadcom said that distributors and equipment makers account for most of its chip sales, and that it builds custom accelerators to each customer's specifications ³⁸. Its AI semiconductor revenue reached \$16.7 billion in the third quarter of fiscal 2026, up 221 percent from a year earlier ⁶. Marvell, another US firm, also builds custom AI chips to each customer's specifications and sells Ethernet switch chips, its annual report for fiscal 2026 says ³⁹. Marvell's five-year agreement with Amazon Web Services, announced in December 2024, covers custom AI chips alongside the optical and networking parts that go with them ⁴⁰.

Taiwanese firms do much of the physical assembly. Foxconn booked NT\$8.1 trillion of revenue in 2025, up 18 percent, and cloud and networking was its largest product category in the fourth quarter ⁷. In the same results release, its chairman called Foxconn the world's largest AI server provider ⁷.

Some assembly is moving to the United States. In April 2025 Nvidia said it was building supercomputer manufacturing plants in Texas, with Foxconn in Houston and Wistron in Dallas ⁴¹. It expected mass production to ramp up in the following 12 to 15 months ⁴¹. In August 2026 Nvidia said Wistron's new plant in Fort Worth, Texas, near Dallas, was producing its GB300 superchips and Foxconn was building its factory in Houston ¹⁶.

The chokepoint

In October 2025 Nvidia said more than 50 makers of systems and components were gearing up for the Vera Rubin NVL72, whose design follows its MGX rack architecture ¹⁸. In its filing for the quarter to July 2026, Nvidia said the complexity of producing its data center systems had caused production delays ²⁹.

38 Broadcom Inc., Form 10-K annual report for the period ended 2025-11-02 (10-K). U.S. Securities and Exchange Commission (filing by Broadcom Inc.) · 18 December 2025 [\[A\]](https://www.sec.gov/Archives/edgar/data/1730168/000173016825000121/avgo-20251102.htm) www.sec.gov/Archives/edgar/data/1730168/000173016825000121/avgo-20251102.htm

6 Broadcom Inc., Form 8-K current report for the period ended 2026-09-02 (8-K). U.S. Securities and Exchange Commission (filing by Broadcom Inc.) · 2 September 2026 [\[A\]](https://www.sec.gov/Archives/edgar/data/1730168/000173016826000076/avgo-08022026x8kxex99.htm) www.sec.gov/Archives/edgar/data/1730168/000173016826000076/avgo-08022026x8kxex99.htm

39 Marvell Technology, Inc., Form 10-K annual report for the period ended 2026-01-31 (10-K). U.S. Securities and Exchange Commission (filing by Marvell Technology, Inc.) · 11 March 2026 [\[A\]](https://www.sec.gov/Archives/edgar/data/1835632/000183563226000011/mrvl-20260131.htm) www.sec.gov/Archives/edgar/data/1835632/000183563226000011/mrvl-20260131.htm

40 Marvell Technology, Inc., Form 8-K current report for the period ended 2024-12-02 (8-K). U.S. Securities and Exchange Commission (filing by Marvell Technology, Inc.) · 2 December 2024 [\[A\]](https://www.sec.gov/Archives/edgar/data/1835632/000183563224000193/final2024_12x02xmarvell-.htm) www.sec.gov/Archives/edgar/data/1835632/000183563224000193/final2024_12x02xmarvell-.htm

7 Hon Hai Technology Group (Foxconn) Announces FY2025 & 4Q25 Financial Results. Hon Hai Precision Industry (Foxconn) · 16 March 2026 [\[A\]](https://www.honhai.com/en-us/press-center/press-releases/latest-news/1978) www.honhai.com/en-us/press-center/press-releases/latest-news/1978

41 NVIDIA to Manufacture American-Made AI Supercomputers in US for First Time. Nvidia · 14 April 2025 [\[A\]](https://blogs.nvidia.com/blog/nvidia-manufacture-american-made-ai-supercomputers-us/) blogs.nvidia.com/blog/nvidia-manufacture-american-made-ai-supercomputers-us/

16 NVIDIA and Partners Build in America, for America. Nvidia · 5 August 2026 [\[A\]](https://blogs.nvidia.com/blog/nvidia-and-partners-build-in-america-for-america/) blogs.nvidia.com/blog/nvidia-and-partners-build-in-america-for-america/


18 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

29 NVIDIA CORP, Form 10-Q quarterly report for the period ended 2026-07-26 (10-Q). U.S. Securities and Exchange Commission (filing by NVIDIA CORP) · 26 August 2026 [\[A\]](https://www.sec.gov/Archives/edgar/data/1045810/000104581026000075/nvda-20260726.htm) www.sec.gov/Archives/edgar/data/1045810/000104581026000075/nvda-20260726.htm

Four components that go into the rack take longer to replace.

- **Indium phosphide laser capacity.** In March 2026 Nvidia agreed to invest \$2 billion in Coherent, alongside a multibillion-dollar purchase commitment and future access and capacity rights for Coherent's advanced laser and optical networking products ⁹.
- **Indium phosphide substrates.** China added indium phosphide substrates, the wafers the lasers are grown on, to its export control list on 4 February 2025 ¹⁴. AXT, a US-listed maker of these wafers, does all its manufacturing in China ¹⁴. In its filing for the quarter to June 2026, AXT called the export permits the most significant challenge it faces ¹⁴. It said it could not estimate when it would resume shipping the wafers to the United States ¹⁴. The same filing names Sumitomo and JX as AXT's competitors in indium phosphide ¹⁴.
- **NVSwitch.** Nvidia's annual report for fiscal 2026 lists NVLink interconnects and switches among its own networking products ¹. UALink is an open standard designed for accelerator communications within one scale-up group ¹⁰. Since May 2025 Nvidia's NVLink Fusion program has let other firms' custom chips use NVLink, including CPUs that Fujitsu and Qualcomm plan to connect to Nvidia GPUs ¹².
- **High-layer-count circuit boards.** TTM Technologies, an American circuit board maker, says its circuit board competitors are mostly based in China and Taiwan ⁴². Nitto Boseki, a Japanese glass fiber maker, makes NE-glass, a glass cloth that absorbs less signal, for the boards in data center switches and routers ⁴³. See Substrates and PCBs.

Optical module assembly can move, as InnoLight's high-volume plant in Thailand, opened in 2019, shows ⁸.



9 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

14 AXT INC, Form 10-Q quarterly report for the period ended 2026-06-30. U.S. Securities and Exchange Commission (filing by AXT INC) · 13 August 2026 [\[A\]](https://www.sec.gov/Archives/edgar/data/1051627/000143774926027677/axti20260630_10q.htm) www.sec.gov/Archives/edgar/data/1051627/000143774926027677/axti20260630_10q.htm

1 NVIDIA CORP, Form 10-K annual report for the period ended 2026-01-25 (10-K). U.S. Securities and Exchange Commission (filing by NVIDIA CORP) · 25 February 2026 [\[A\]](https://www.sec.gov/Archives/edgar/data/1045810/000104581026000021/nvda-20260125.htm) www.sec.gov/Archives/edgar/data/1045810/000104581026000021/nvda-20260125.htm

10 UALink™ 200G 1.0 Specification Overview. UALink Consortium · 15 May 2025 [\[A\]](https://ualinkconsortium.org/blog/ualink-200g-1-0-specification-overview-802/) ualinkconsortium.org/blog/ualink-200g-1-0-specification-overview-802/

12 NVIDIA Unveils NVLink Fusion for Industry to Build Semi-Custom AI Infrastructure With NVIDIA Partner Ecosystem. Nvidia · 18 May 2025 [\[A\]](https://nvidianews.nvidia.com/news/nvidia-nvlink-fusion-semi-custom-ai-infrastructure-partner-ecosystem) nvidianews.nvidia.com/news/nvidia-nvlink-fusion-semi-custom-ai-infrastructure-partner-ecosystem

42 TTM TECHNOLOGIES INC, Form 10-K annual report for the period ended 2025-12-29 (10-K). U.S. Securities and Exchange Commission (filing by TTM TECHNOLOGIES INC) · 17 February 2026 [\[A\]](https://www.sec.gov/Archives/edgar/data/1116942/000119312526051976/ttmi-20251229.htm) www.sec.gov/Archives/edgar/data/1116942/000119312526051976/ttmi-20251229.htm

43 Nittobo Group Integrated Report 2025. Nitto Boseki (Nittobo) · 16 September 2025 · p. 7 [\[A\]](https://www.nittobo.co.jp/eng/ir/pdf/Nittobo_Integrated-Report_2025.pdf#page=7) www.nittobo.co.jp/eng/ir/pdf/Nittobo_Integrated-Report_2025.pdf#page=7

8 InnoLight Technology, company overview page. InnoLight Technology · Accessed 9 September 2026 [\[A\]](https://www.innolight.com/about) www.innolight.com/about

FOUR THINGS TO REMEMBER

- An AI rack wires 72 chips together to run as one computer.
- Almost every link longer than 7 meters carries its signal as light, because copper fades over longer runs.
- Nvidia's NVLink switch chips and the lasers are the hardest parts to replace. In October 2025 more than 50 makers of systems and components were gearing up to build Nvidia's next rack.
- China produced about 70 percent of the world's indium in 2025 and has required export permits for indium phosphide substrates since February 2025.

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

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. "Systems and Networking." Chip Supply Chain: How AI Chips Are Made. Edition of 20 October 2026. <https://chipsupplychain.org/systems-and-networking/>

```
@misc{chipsupplychain2026systemsandnetworking,
  title      = {Systems and Networking},
  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/systems-and-networking/}
}
```