



xETFs

NECK

xETFs AI Bottlenecks ETF

AI's Critical Constraints.
One Ticker.

September 2026

© 2026 WallStreetX ETFs, Inc.

Disclaimer

Carefully consider the Funds' investment objectives, risk factors, charges and expenses before investing. This and additional information can be found in the Funds' Prospectus and Summary Prospectus, which may be obtained by visiting www.xETFs.com/NECK. Read the Prospectus and Summary Prospectus carefully before investing.

Market figures are third-party estimates, and actual outcomes may differ from estimates. Past performance is not a guarantee of future results. xETFs cannot guarantee the accuracy of third-party information, including estimates, and does not represent that such information is accurate or complete.

Exchange Traded Concepts, LLC serves as the investment adviser. WallStreetX ETFs, Inc. dba xETFs serves as the sub-adviser. The Funds are distributed by Foreside Fund Services, LLC., which is not affiliated with xETFs, Exchange Traded Concepts, LLC, or any of its affiliates.

Investing involves risk, including possible loss of principal. To the extent the Fund's investments are concentrated in or have significant exposure to a particular issuer, industry or group of industries, or asset class, the Fund may be more vulnerable to adverse events affecting such issuer, industry or group of industries, or asset class than if the Fund's investments were more broadly diversified. Issuer-specific events, including changes in the financial condition of an issuer, can have a negative impact on the value of the Fund.

A new or smaller fund is subject to the risk that its performance may not represent how the fund is expected to or may perform in the long term. In addition, new funds have limited operating histories for investors to evaluate, and new and smaller funds may not attract sufficient assets to achieve investment and trading efficiencies.

Shares are bought and sold at market price (closing price) not net asset value (NAV) and are not individually redeemed from the Fund. Market price returns are based on the midpoint of the bid/ask spread at 4:00pm Eastern Time (when NAV is normally determined) and do not represent the return you would receive if you traded at other times. Brokerage commissions will reduce returns.

AI Bottleneck Companies Risk. The Fund invests in AI Bottleneck Companies that provide products, technologies or services critical to the development, deployment or scaling of AI, including memory; optics, photonics and networking; power and infrastructure; semiconductors and compute; and other AI infrastructure bottlenecks. These companies may be highly dependent on the continued growth of AI-related spending, infrastructure demand and commercialization. If AI adoption, investment or monetization is slower than expected, or demand for AI-related infrastructure is overestimated, the securities of these companies may experience significant declines. Companies exposed to AI bottlenecks may be subject to rapid technological change, intense competition, short product cycles, supply and demand imbalances, component shortages, manufacturing constraints, customer concentration, and significant development and capital requirements. Their businesses may also be adversely affected by export controls, trade restrictions, intellectual property disputes, cybersecurity incidents, energy constraints, geopolitical tensions, and disruptions in key technology supply chains. A company identified as addressing an AI bottleneck may lose that position as technology evolves, competitors enter the market, production capacity expands, or alternative technologies reduce the need for its products or services. There can be no assurance that any identified bottleneck will result in increased revenues, margins, pricing power, or stock price appreciation for those companies. Bottlenecks may prove temporary, may be resolved more quickly than expected, or may shift to other parts of the AI stack as technology evolves, competitors enter the market, production capacity expands, or alternative technologies reduce the need for a company's products or services.

Asia Risk. The Fund expects to invest significantly in the securities of Asian issuers. As such, the Fund is subject to certain risks specifically associated with investments in the securities of Asian issuers. Many Asian economies have experienced rapid growth and industrialization, and there is no assurance that this growth rate will be maintained. Some Asian economies are highly dependent on trade, and economic conditions in other countries within and outside Asia can impact these economies. Certain of these economies may be adversely affected by trade or policy disputes with its major trade partners. There is also a high concentration of market capitalization and trading volume in a small number of issuers representing a limited number of industries, as well as a high concentration of investors and financial intermediaries. Certain Asian countries have experienced and may in the future experience expropriation and nationalization of assets, confiscatory taxation, currency manipulation, political instability, armed conflict and social instability as a result of religious, ethnic, socio-economic and/or political unrest. In particular, escalated tensions involving North Korea and any outbreak of hostilities involving North Korea could have a severe adverse effect on Asian economies. Governments of certain Asian countries have exercised, and continue to exercise, substantial influence over many aspects of the private sector. In certain cases, the government owns or controls many companies, including the largest in the country. Accordingly, government actions could have a significant effect on the issuers of the Fund's securities or on economic conditions generally. Recent developments in relations between the U.S. and China have heightened concerns of increased tariffs and restrictions on trade between the two countries. An increase in tariffs or trade restrictions, or even the threat of such developments, could lead to a significant reduction in international trade, which could have a negative impact on the economy of Asian countries and a commensurately negative impact on the Fund.

Disclaimer Continued

Concentration Risk. The Fund is susceptible to an increased risk of loss, including losses due to adverse events that affect the Fund's investments more than the market as a whole, to the extent that the Fund's investments are concentrated in the securities and/or other assets of a particular issuer or issuers, country, group of countries, region, market, industry, group of industries, sector, market segment or asset class.

Depository Receipts Risk. The Fund may invest in depository receipts. Investment in ADRs and GDRs may be less liquid than the underlying shares in their primary trading market and GDRs, many of which are issued by companies in emerging markets, may be more volatile and less liquid than depository receipts issued by companies in more developed markets.

Derivatives Risk. The use of derivative instruments involves risks different from, or possibly greater than, the risks associated with investing directly in securities and other traditional investments. These risks include: (i) the risk that the counterparty to a derivative transaction may not fulfill its contractual obligations; (ii) risk of mispricing or improper valuation; and (iii) the risk that changes in the value of the derivative may not correlate perfectly with the underlying asset. Derivative prices are highly volatile and may fluctuate substantially during a short period of time.

Foreign Securities Risk. Investments in non-U.S. securities involve certain risks that may not be present with investments in U.S. securities. For example, investments in non-U.S. securities may be subject to risk of loss due to foreign currency fluctuations or to expropriation, nationalization or adverse political or economic developments.

Geographic Investment Risk. To the extent the Fund invests a significant portion of its assets in the securities of companies of a single country or region, it is more likely to be impacted by events or conditions affecting that country or region.

Semiconductor Companies Risk. The semiconductor industry is characterized by rapid technological change and product obsolescence, cyclical market patterns, price erosion, periods of over-capacity and production shortages, variations in manufacturing costs and yields, and significant expenditures for capital equipment and product development.

Small Capitalization Companies Risk. Small capitalization companies may be more vulnerable to adverse general market or economic developments, and their securities may be less liquid and may experience greater price volatility than large and mid-capitalization companies as a result of several factors, including limited trading volumes, fewer products or financial resources, management inexperience and less publicly available information. Accordingly, such companies are generally subject to greater market risk than large and mid-capitalization companies.

AI Is No Longer a Sector. It Is Embedded in the Market.

The largest companies are essential to spending on Artificial Intelligence ("AI") infrastructure.



The ten largest U.S. companies are all technology related

The ten largest U.S. companies by market capitalization are all technology businesses: the Magnificent Seven plus SpaceX (SPCX), Broadcom (AVGO), and Micron (MU).¹ Each one is related to the AI build-out: semiconductors, cloud infrastructure, foundation models, or edge devices. A fund invested in the U.S. market holds substantial exposure to the AI theme.



The AI buildout is large enough to move GDP

Federal Reserve staff estimate data-center investment added roughly 0.41 percentage points to nominal U.S. GDP growth in 2025 and expect the 2026 contribution to be from 0.75 to 1.11 points.⁴



The AI buildout is already in reported earnings

NVIDIA reported \$89.0bn of data-center revenue in a single quarter, up 117% year over year.⁵ This revenue requires suppliers of components - memory, optics, power and packaging.

~53%

of the S&P 500 is tech and AI-linked

Information Technology is 37.9% of the index, Communication Services another 9.5%,² and Amazon and Tesla add 5.4% more.³

\$89.0bn

NVIDIA data-center revenue

In one quarter (fiscal Q2 2027, ended July 2026), up 117% year over year, on total revenue of \$96.2bn.⁵

Capital is Committed to Meet the Demand for AI

The buyers of AI hardware are funding the physical buildout directly, and NVIDIA is funding some of its own suppliers.

\$500bn+

NVIDIA-partnered financing platforms

NVIDIA announced strategic partnerships with Apollo, BlackRock, Blackstone, Brookfield, Goldman Sachs and KKR, to establish independent compute financing platforms, for over \$500bn of capital for AI infrastructure buildout.¹

₩1,100tn

SK hynix long-term Korean domestic investments

Pledged in June 2026 for an AI memory belt linking Yongin, Cheongju and a new southwest cluster, with completion targeted by 2033.²

~\$700bn+

2026 hyperscaler capex

Committed 2026 capital expenditure across the five largest cloud platforms, roughly double 2025 levels.³

\$60–64bn

TSMC 2026 capital budget

Up about 50% on 2025 actual capex of \$40.9bn. The foundry that builds the accelerators is spending to keep pace with them.⁴



NVIDIA is now investing in its own suppliers

In March 2026, NVIDIA announced plans to make strategic investments of \$ 2bn each into photonics companies Coherent and Lumentum, each alongside a multibillion-dollar purchase commitment and future access to capacity. NVIDIA also announced plan to invest \$2bn in Marvell Technology, a semiconductor company.⁵

What It Takes to Run an AI Model

Every AI token generated is read out of memory, moved across a network, paid for in electricity, and only then computed.

1

The model has to sit in memory

An AI model is billions of stored numbers. To answer a single query, the chip has to read all of them. Those numbers live in high-bandwidth memory bolted directly onto the processor, and every new chip generation carries more of this memory.

80GB on a 2023 H100, 288GB on NVIDIA's Rubin.¹

2

The chips connect and communicate

One chip is nowhere near enough, so thousands are wired together and made to behave like a single computer. Inside a rack that wiring is copper. Past the rack, it has to become light.

A Vera Rubin rack links 72 chips holding 20.7 TB of memory.¹

3

The data center has to power it

Serving a model is not a burst of activity. It is a constant load, every second of every day, that has to be generated, delivered to the building, and then cooled once it turns into heat.

Power is consumed regardless of whether the model is busy.

4

And only then does it compute

With the weights in memory, the chips linked and the power flowing, the cluster can finally do what it was built for: train the next model, or serve tokens to millions of users at once. Everything upstream exists to keep this one step fed.

One Vera Rubin rack delivers 3.6 exaFLOPS of FP4 inference.¹

Demand Compounds. Supply Is Physical.

Across multiple layers of AI inputs, indications of constrained capacity.



MEMORY

“We still forecast that customer demand will remain higher than our supply capacity even beyond 2030.”

Kwak Noh-Jung, CEO, SK hynix, July 2026. He separately flagged 2027 as the worst memory supply year on record.¹



OPTICS & NETWORKING

800G optical transceivers supply is expected to fall 40% - 60% short of demand through 2027, and 1.6T transceivers shortfall of 30% - 40% is expected through 2029.

McKinsey & Co., Opportunities in networking optics: Boosting supply for data centers, June 2025.²



POWER & INFRASTRUCTURE

“Limited power, land and infrastructure are slowing development and keeping vacancy near zero in some key U.S. markets.”

Head of CBRE Data Center Research, Americas, June 2026. Global data-center vacancy was 6.7% in Q1 2026, with some key markets much lower.³



SEMICONDUCTORS & COMPUTE

“Our packaging capacity is so tight that now it’s limiting my customers’ growth.”

C.C. Wei, CEO, TSMC, July 2026.⁴

Multiple Physical Constraints, One Portfolio

NECK will focus on four key areas of constraints for 80%+ of the portfolio and will also invest in potential future bottlenecks



MEMORY

Holds the model while it runs



OPTICS, PHOTONICS, & NETWORKING

Connects the chips to each other



POWER & INFRASTRUCTURE

Generates, delivers and cools the electricity



SEMICONDUCTORS & COMPUTE

Designs, builds and packages the chips



POTENTIAL FUTURE BOTTLENECKS

Potential new areas of constraints



NECK will rebalance its portfolio periodically, to reflect changes in bottlenecks over time

Relief at one layer pushes pressure to the next. More wafers need more packaging. More packaging needs more memory. More memory needs more power, and more power needs more transformers and more turbine slots.

What Makes Something a Bottleneck?

Four assessments. Being critical is not enough.



Demand exceeds supply

Sold-out capacity, order backlogs, allocation, multi-year lead times and lengthening qualification cycles.



Pricing power

Tight supply supports higher prices, and those gains translate into stronger margins.



Durable moats


Process IP, scale, customer qualification cycles and capital barriers that take years, not quarters, for a competitor to replicate.



Financial materiality

The constrained product must be large enough within the business to meaningfully impact revenue, earnings, or margins.

The test we apply hardest: a critical component is not automatically a good investment. A part can be indispensable and still be commoditized, second-sourced, or priced by a customer with more leverage than the supplier. We look for the places where scarcity reaches the income statement.

The background of the slide features a dark blue color scheme. On the right side, there is a perspective view of several rows of server racks. Overlaid on the entire background are numerous thin, white, wavy lines that flow from the bottom left towards the top right, creating a sense of movement and data flow.

PART TWO

Get to Know NECK

An actively managed ETF holding 20 companies at launch across the layers where AI demand meets a physical limit.

Introducing NECK

xETFs AI Bottlenecks ETF

Fund Snapshot

Fund ticker	NECK
Investment objective	Capital appreciation
Structure	U.S.-listed, actively managed ETF
Listing exchange	NYSE
Holdings	20 at launch; 15 to 25 expected
Strategy	Seeks to invest in companies positioned at critical points of constraint across the artificial intelligence (AI) stack, meaning companies whose products or capacity are constrained relative to demand, such as being sold out, carrying a backlog, or having capacity already committed, or whose position in the AI supply chain is difficult to replace (“AI Bottleneck Companies”).
Tax reporting	1940 Act ETF; no K-1s. Dividends from net investment income paid annually
Adviser / Sub-adviser	Exchange Traded Concepts, LLC / WallStreetX ETFs, Inc. (xETFs)
Inception Date	September 23, 2026
Total Annual Fund Operating Expenses	0.75%



What NECK Offers

- Focused exposure to key constraint layers: memory; optics, photonics & networking; power & infrastructure; semiconductors & compute
- Companies where scarcity is visible in backlogs, lead times, pricing and margins
- Periodically rebalanced to reflect changes in constraints



What It Is Not

- Not a mega-cap AI fund: no hyperscalers and no graphics processing unit (GPU) designers at the center of the portfolio
- Not a broad semiconductor fund: power, cooling and optics sit alongside silicon

Multiple Bottlenecks. One Ticker.

NECK seeks exposure across the AI infrastructure stack. Areas of portfolio allocation at launch include:

% of fund by holding; grouping totals below

MEMORY	OPTICS, PHOTONICS & NETWORKING	POWER & INFRASTRUCTURE	SEMICONDUCTORS & COMPUTE	FUTURE BOTTLENECKS
HBM, DRAM & NAND	Lasers, Transceivers, Digital Signal Processors (DSPs) & Switching	Turbines, Switchgear, Thermal & On-Site Power	Lithography, Foundry, Packaging & Capacity	Potential Architectures That Route Around the Constraint
SK hynix10.0%	Coherent8.0%	Vertiv Holdings4.5%	ASML Holding6.5%	Cerebras Systems3.5%
Samsung Electronics8.5%	Lumentum8.0%	Eaton Corp4.5%	TSMC4.5%	Oklo2.5%
Micron Technology7.0%	Marvell Technology4.5%	GE Vernova4.5%	Broadcom4.0%	Sivers Semiconductors2.5%
Kioxia Holdings4.5%	Arista Networks4.5%	Bloom Energy4.0%	CoreWeave2.5%	Everspin1.5%
30.0%	25.0%	17.5%	17.5%	10.0%

Read it as a chain, not a basket. Memory feeds the accelerator, optics connects the accelerators, power runs them, silicon builds them, and Future Bottlenecks holds the companies trying to make those links unnecessary.

The Portfolio at Launch: 20 Names

Grouped by constraint, then ordered by weight within each.

Company	Ticker	Weight	Why It Is in NECK
MEMORY		30.0%	
SK hynix	000660 KS	10.0%	HBM and dynamic random-access memory (DRAM) leader; NVIDIA's lead memory partner
Samsung Electronics	005930 KS	8.5%	The only firm running DRAM, NAND flash memory (NAND) and a leading-edge foundry at scale
Micron Technology	MU US	7.0%	U.S. memory pure-play with multi-year customer agreements
Kioxia Holdings	285A JP	4.5%	100% NAND pure-play as AI inference turns storage into a constraint
OPTICS, PHOTONICS & NETWORKING		25.0%	
Coherent Corp	COHR US	8.0%	Indium phosphide substrates, laser die and transceivers, vertically integrated
Lumentum Holdings	LITE US	8.0%	Laser chips, cloud transceiver modules and ultra-high-power CPO lasers
Marvell Technology	MRVL US	4.5%	Optical DSPs, 1.6T interconnect and custom AI silicon
Arista Networks	ANET US	4.5%	Data-center Ethernet switching, the fabric that clusters GPUs
POWER & INFRASTRUCTURE		17.5%	
Vertiv Holdings	VRT US	4.5%	Power distribution and thermal management for high-density AI racks
Eaton Corp	ETN US	4.5%	Electrical distribution and switchgear; the longest lead-time gear in the chain
GE Vernova	GEV US	4.5%	Gas turbines and grid equipment, the gate on time-to-power
Bloom Energy	BE US	4.0%	On-site fuel cells that bypass the grid interconnection queue
SEMICONDUCTORS & COMPUTE		17.5%	
ASML Holding	ASML US	6.5%	Sole supplier of EUV lithography
TSMC	TSM US	4.5%	Leading-edge foundry and CoWoS advanced packaging
Broadcom	AVGO US	4.0%	Custom AI accelerators plus the switch and optical silicon around them
CoreWeave	CRWV US	2.5%	GPU cloud; contracted accelerator capacity sold forward
FUTURE BOTTLENECKS		10.0%	
Cerebras Systems	CBRS US	3.5%	Wafer-scale AI chips that remove the off-chip interconnect
Oklo Inc	OKLO US	2.5%	Small modular reactors delivering firm baseload behind the meter
Sivers Semiconductors	SIVE SS	2.5%	Merchant indium phosphide laser die for silicon photonics
Everspin Technologies	MRAM US	1.5%	MRAM persistent memory, an alternative to the DRAM/NAND trade-off

Memory: The Wall Every Model Hits

The binding constraint on AI performance has moved from processing speed to memory bandwidth and capacity.



A fast chip with slow memory is a slow chip

Modern accelerators can perform arithmetic far faster than memory can feed them data, so a chip spends much of its time idle. Every token a model generates requires pulling its weights back out of memory, so bandwidth and capacity, not raw processing power, set how large a model can be, how much context it can hold, and how many users a single machine can serve.

3.6x

more memory per chip

From 80GB on a 2023 H100 to 288GB on Rubin.¹



Every generation needs more of it

Each new accelerator carries more memory than the last, and every stack has to be made by one of three suppliers. Stacking it is a separate manufacturing step with its own capacity limit.



Memory for AI eats ordinary memory

The chips inside high-bandwidth memory (HBM) are made on the same lines as normal dynamic random-access memory (DRAM). A wafer sent to one is a wafer not available for the other, which is how an AI shortage becomes an everything shortage.



And there are only three suppliers

SK hynix held about 58% share of high-bandwidth memory (HBM) revenue in Q1 2026, with Samsung and Micron at roughly 21% each. Three companies supply one of the critical inputs in the AI stack.²

Memory is a Scarce Resource

Memory scarcity is already visible in reported earnings, not just in order books.

557% YoY

SK hynix Q2 2026 operating income change vs Q2 2025

On YOY Revenue increase of 257%.¹

85%

Micron gross margin

Fiscal Q3 2026 GAAP gross margin, on revenue of \$41.5bn. Q4 guidance was \$50.0bn revenue at roughly 86% gross margin.²

Over 60%

Kioxia annual cap ex spending vs 2025

Planned ¥470bn average spending each year for 2026 –2028 versus ¥284bn in FY2025, in response to expansion of NAND flash memory (NAND) market driven by AI and datacenter demand.³

\$38.8bn

2026 DRAM equipment sales

SEMI's forecast for DRAM equipment in 2026, with projected increase to \$56.9bn by 2028.⁴

“We still forecast that customer demand will remain higher than our supply capacity even beyond 2030.”

Kwak Noh-Jung, Chief Executive Officer, SK hynix, July 2026. In the same remarks he flagged 2027 as the worst memory supply year on record.⁵

Memory: Initial Target 30% of NECK

Three dynamic random-access memory (DRAM) makers and the NAND flash memory (NAND) pure-play, chosen for how directly memory scarcity reaches earnings.

10.0%

SK hynix

000660 SK · HBM & DRAM

The purest expression of the thesis. Roughly 58% of global HBM revenue and NVIDIA's lead memory partner, with memory essentially 100% of revenue, so tightness lands in earnings undiluted. Q2 2026 operating margin was 76% on revenue up 257% year over year.^{1, 2}

8.5%

Samsung Electronics

005930 KS · Memory, foundry & devices

The only company running DRAM, NAND and a leading-edge foundry at scale, and the largest DRAM supplier by revenue at 38% share in Q1 2026.² Its broader business mix means less direct memory exposure than a pure-play supplier, but provides diversification alongside meaningful HBM4 upside.²

7.0%

Micron Technology

MU US · DRAM, HBM & NAND

The U.S. memory pure-play, and the clearest evidence that scarcity is being priced. Fiscal Q3 2026 revenue of \$41.5bn at an 80.4% operating margin, with Q4 guidance of revenues at \$50.0bn. Multi-year customer agreements convert a shortage into contracted visibility.³

4.5%

Kioxia Holdings

285A JP · NAND flash

A 100% NAND pure-play at the moment AI inference turns storage into a bottleneck. Kioxia projects data-center flash demand at a 46% CAGR to 2028 and inference workloads at 86%, and is lifting capex about 60% versus 2025 levels, to roughly ¥470bn a year to meet it.⁴

Optics, Photonics & Networking: Why AI Clusters Need It

Past a few meters, copper cannot carry the data rates AI networking now runs at. Every connection beyond the rack becomes optical.

At 800G and 1.6T per lane, copper runs out at about three meters


COPPER

usable

signal degrades: too much loss and noise to recover the data

OPTICAL

usable across the whole building, at full rate



So the count scales with the cluster

A cluster is not one machine, it is thousands wired together. Every link that leaves a rack needs a transceiver at each end, so the number of optical parts grows with the size of the cluster, not with the number of buildings.



Which repriced the whole market

Optical transceiver and related products sales reached \$26.8bn in 2025, up 74% in a year.¹ AI – suitable transceivers have projected supply shortages: 800G optical transceivers supply is expected to be 40% - 60% short of demand through 2027, and 1.6T transceivers shortfall of 30% - 40% is expected through 2029.²



And it is now moving onto the chip

Co-packaged optics puts the light source beside the switch chip instead of in a plug-in module. NVIDIA already ships this in Spectrum-X Photonics, which moves optics from an accessory into the semiconductor supply chain.³

Inside the Optical Supply Chain Constraint

Constraints exist at multiple points in the chain, and no single stage relieves them all



Indium phosphide wafer

The specialist crystal a laser has to be grown on. Not silicon, and not made on silicon lines.

Coherent



Laser die

The light source itself. A handful of qualified suppliers worldwide, and capacity that takes years to add.

Coherent · Lumentum



Optical engine

Laser, modulator and detector packaged together with the chip that encodes the signal.

Coherent · Lumentum
Marvell



Transceiver module

The finished part that plugs into a switch, or is bonded beside the switch chip.

Coherent · Lumentum



Switch & fabric

The network that makes thousands of accelerators behave as a single machine.

Arista · Marvell



Why upstream scarcity matters

Upstream, Indium Phosphide (InP) wafer and laser capacity is limited by specialized fabs, tools, yields and customer qualification, so new supply can take years. Downstream, Digital Signal Processors (DSPs), switch silicon and modules are constrained by design cycles, qualification and silicon supply instead.



Different constraints, different exposures

The supply chain is constrained at more than one point, and the holdings sit across all of them. Upstream in wafer and laser capacity, and downstream in the silicon and switches the module plugs into.

Optics, Photonics & Networking: Initial Target 25% of NECK

From the laser die to the fabric, four companies that carry the signal between accelerators.

8.0%

Coherent Corp

COHR US · Lasers & optical components

Vertically integrated from indium phosphide substrate through to the finished transceiver. Q4 FY2026 revenue of \$2.05bn, up 34% year over year, with gross margin up 277 basis points as mix shifts to datacom.¹

8.0%

Lumentum Holdings

LITE US · Transceivers & laser chips

The steepest revenue inflection in the grouping: FY2026 revenue of \$3.01bn, up 83%, with Q4 alone more than doubling year over year and non-GAAP gross margin rising from 34.7% to 46.0%.²

4.5%

Marvell Technology

MRVL US · Optical DSPs & custom silicon

The silicon between the optics and the compute: optical DSPs, 1.6T interconnect and custom accelerators. Data center is 76% of revenue; Q1 2027 results reflected strength in the data center end market.³

4.5%

Arista Networks

ANET US · Data-center Ethernet switching

Supplies the AI fabric/ switching that turns thousands of GPUs into one machine. Q2 2026 was its first \$3bn revenue quarter, up 38% year over year. Full year revenue guidance for 2026 was raised to \$12.6 billion, representing 40% annual growth.⁴

Power & Infrastructure: One of the Hardest Resources to Buy

A finished data center with no connection to the grid is useless. The scarce asset is a delivery date.

2030-2031

Order pipeline for new gas-turbines

GE Vernova reported its orders were full through 2029 and was taking orders through 2031 deliveries.¹

~945 TWh

Estimated global data-center use by 2030²

Up from roughly 485 TWh in 2025, with use by AI-focused data centers roughly tripling over the same period.²

9-17%

Estimates for data centers' share of U.S. power, per research

EPRI's 2030 scenario range for data centers as a share of all U.S. electricity, up from roughly 4-5% today. Lawrence Berkeley National Laboratory's central case is 11.8%.³

0.3%

Northern Virginia vacancy

The world's largest data-center market is effectively full. Global vacancy was 6.7% in Q1 2026, down from 8.3% a year earlier.⁴



The queue is the constraint, not the capital

Interconnection queues, transformer and cable lead times, and gas-turbine slot availability are measured in years. Money is abundant; a delivery date is not. GE Vernova held 100-116 GW of gas backlog including slot reservations in mid-2026.⁵



Every megawatt runs through the same suppliers

NECK focuses on companies whose lead times decide when electricity arrives.

What Actually Delivers the Power to AI

A megawatt reaches a chip through multiple methods, and NECK seeks exposure to these.

Held in NECK



Gas turbines

Dispatchable generation at utility scale. Slots are booked years out: As of Q2 2026, 116 Gigawatts of gas backlog including slot reservations at GE Vernova, with 125 Gigawatts expect by yearend 2026.¹

GE Vernova



On-site fuel cells

Behind-the-meter generation that skips the interconnection queue entirely. In 2025, Bloom delivered a fully operational, modular fuel cell system in 55 days against a standard 90-day timeline.²

Bloom Energy



Grid & electrical equipment

Transformers, switchgear, busway and UPS: the gear between the grid and the rack, and the longest lead-time items in the chain. Eaton's data-center backlog as of Q2 2026 was 307 gigawatts, or 15 years of backlog at 2025 build rates.³

Eaton · GE Vernova



Thermal management

Rack densities are climbing toward 100kW and beyond, past the point where air cooling works. Liquid cooling moves from an option to a requirement.⁴

Vertiv



Advanced nuclear (emerging)

Potential new capacity. The US Nuclear Regulatory Commission (NRC) approved Oklo's principal design criteria report for a nuclear powerhouse under construction in Idaho.⁵

Oklo, in Future Bottlenecks

Power & Infrastructure: Initial Target 17.5% of NECK

Four companies standing between a signed data center lease and an AI running chip.

4.5%

Vertiv Holdings

VRT US · Power & thermal systems

A near pure-play on digital infrastructure: power distribution and thermal management for high-density rack needs. Q2 2026 sales of \$3.27bn, up 24%, with full-year organic growth guided to 30-32% and Q3 organic growth to 34-36%.¹

4.5%

Eaton Corp

ETN US · Electrical distribution & switchgear

One of a handful of scaled suppliers in power distribution and switchgear, where lead times remain elevated. Electrical Americas twelve-month rolling orders were up 41% organically with backlog up 33%; data-center orders rose roughly 85%.²

4.5%

GE Vernova

GEV US · Gas turbines & grid equipment

The gate on time-to-power, in an effective three-player gas turbine oligopoly. \$176bn total backlog, gas backlog including slot reservations of 100-116 GW, and an electrification equipment backlog of \$40.6bn, up 69% year over year.³

4.0%

Bloom Energy

BE US · On-site fuel-cell power

Offers a solution for power in months. Q2 2026 revenue of \$1.07bn, up 166%, swinging to \$182m of operating income, with an Oracle master agreement for up to 2.8 Gigawatts.⁴

Semiconductors & Compute: AI Runs on Leading-Edge Silicon

One lithography supplier, one foundry with nearly all leading-edge AI volume, and a packaging step that is currently limiting customers.




One supplier for the hardest step

ASML Holding N.V. is the only company in the world that makes EUV (extreme ultraviolet) lithography machines, the tools that print the finest features on a chip. Capacity expands on a multi-year clock: 2027 output is targeted to expand about 30% above 2026's roughly 65 systems.¹



One foundry for the leading edge

Taiwan Semiconductor Manufacturing Company Limited (TSMC) makes the overwhelming majority of leading-edge AI silicon. Capex was raised to \$60-64bn for 2026, spending that will produce output for the long term.²



And packaging may be what runs out first

A finished silicon wafer is not a finished chip. Memory still has to be stacked onto the logic and sealed into one package and requires separate capacity. Today it is the step that constrains before wafers do.²

From a Design to a Working Rack

Four steps between an idea and delivered compute, each controlled by a handful of suppliers.

STEP 1



Design the chip for AI

Broadcom co-designs custom accelerator chips with the hyperscalers, the main alternative to merchant GPUs.

Broadcom

STEP 2



Print the logic

ASML's machines print the processor onto a silicon wafer at TSMC, a foundry at the leading edge.

ASML · TSMC

STEP 3



Join it together

Logic and memory are bonded side by side onto one carrier and sealed into a single package.

TSMC

STEP 4



Build the rack

The finished package becomes a board, then a rack, then a cluster wired with optics, powered and cooled, and then rented out for use.

CoreWeave



Every step is a critical step.

There is a single supplier of EUV lithography, one foundry that captures the overwhelming majority of leading-edge AI volume, and a packaging step that TSMC's chief executive says is "so tight that now it's limiting my customers' growth." Broadcom's AI semiconductor revenue in Q2 FY26 rose 143% year over year as custom accelerators took share.^{1, 2}

Semiconductors & Compute: Initial Target 17.5% of NECK

The tool, the fab, the custom chip and the rented capacity.

6.5%

ASML Holding

ASML US · EUV lithography

There is no second source for EUV. Q2 2026 net sales of €9.3bn with FY2026 guidance raised to €43-45bn, and low-NA EUV capacity planned 30% higher in 2027 with another 30% under study for 2028.¹

4.5%

TSMC

TSM US · Foundry & CoWoS packaging

No credible alternative at the leading edge, and the owner of the packaging step currently limiting customer growth. Q2 2026 revenue of \$40.2bn, up 34%, at a 60.3% operating margin, with high-performance computing at 66% of revenue.²

4.0%

Broadcom

AVGO US · Custom accelerators & network silicon

The main alternative to merchant GPUs, plus the switch and optical silicon around them. Q2 FY2026 AI semiconductor revenue of \$10.8bn, up 143%; guidance of \$16.0bn in Q3 for growth above 200% year over year.³

2.5%

CoreWeave

CRWW US · GPU cloud (neocloud)

Rented scarcity: contracted accelerator capacity sold years forward. Q2 2026 revenue of \$2.58bn, up 112%, on a \$104bn revenue backlog and roughly 3.7 GW of contracted power.⁴

Potential Future Bottlenecks: Cerebras & Oklo

Not every constraint gets relieved by building more of the same thing. Some could get routed around.



Cerebras Systems

CBRS US

3.5%

Routes around: wiring thousands of chips together

Instead of connecting thousands of separate chips, Cerebras builds one enormous chip the size of a dinner plate. Everything stays on a single piece of silicon, so the optics and memory links that exist to move data between chips are simply not needed.

- Q2 2026 revenue of \$180.1m, up 74% year over year, with cloud revenue up 281%¹
- \$25.4bn of remaining performance obligations (backlog of revenue for contracted work), over 90% due to 1 Open AI contract¹
- Risks include but not limited to: few large customers, negative operating margin, and a design that has to win against an entrenched incumbent¹



Oklo Inc

OKLO US

2.5%

Routes around: waiting years for a grid connection

Builds small nuclear reactors designed to sit next to the data center and power it directly. If it works, a developer no longer joins a multi-year queue for transmission capacity, which is currently what decides where data centers can be built at all.

- Potential new capacity. The US Nuclear Regulatory Commission (NRC) approved Oklo's principal design criteria report for a nuclear powerhouse under construction in Idaho.²
- Risks include but not limited to: no revenue yet, and licensing, fuel supply and construction timelines are largely outside the company's control

Potential Future Bottlenecks: Sivers & Everspin

Not every constraint gets relieved by building more of the same thing. Some could get routed around.



Sivers Semiconductors

SIVE SS

2.5%

Routes around: the shortage of laser chips

Makes the tiny lasers that put data onto light. These lasers are a critical input for both today's pluggable optics and next-generation co-packaged optics (CPO). Almost every optics company needs them and very few companies can make them.

- In-house indium phosphide production plus a laser fab in Glasgow¹
- Opportunity pipeline grew 77% year to date to \$799m¹
- Risks include but not limited to: small company with lumpy, program-driven revenue. Q1 2026 sales fell 22% year over year¹



Everspin Technologies


MRAM US

1.5%

Routes around: the trade-off between the two memory types

Today a system either uses fast memory that forgets everything when power is lost, or slow storage that remembers. Everspin makes a third kind that is nearly as fast as the first and keeps its contents like the second.

- The only company supplying this type of memory in volume
- Q2 2026 revenue of \$18.7m, up 42% year over year, at a 53.9% gross margin²
- Risks include but not limited to: today's demand has mostly been for industrial and defense uses. The AI use case is future optionality²



xETFs

AI's Bottlenecks. One Ticker.

NECK · xETFs AI Bottlenecks ETF · NYSE

September 2026

© 2026 WallStreetX ETFs, Inc.