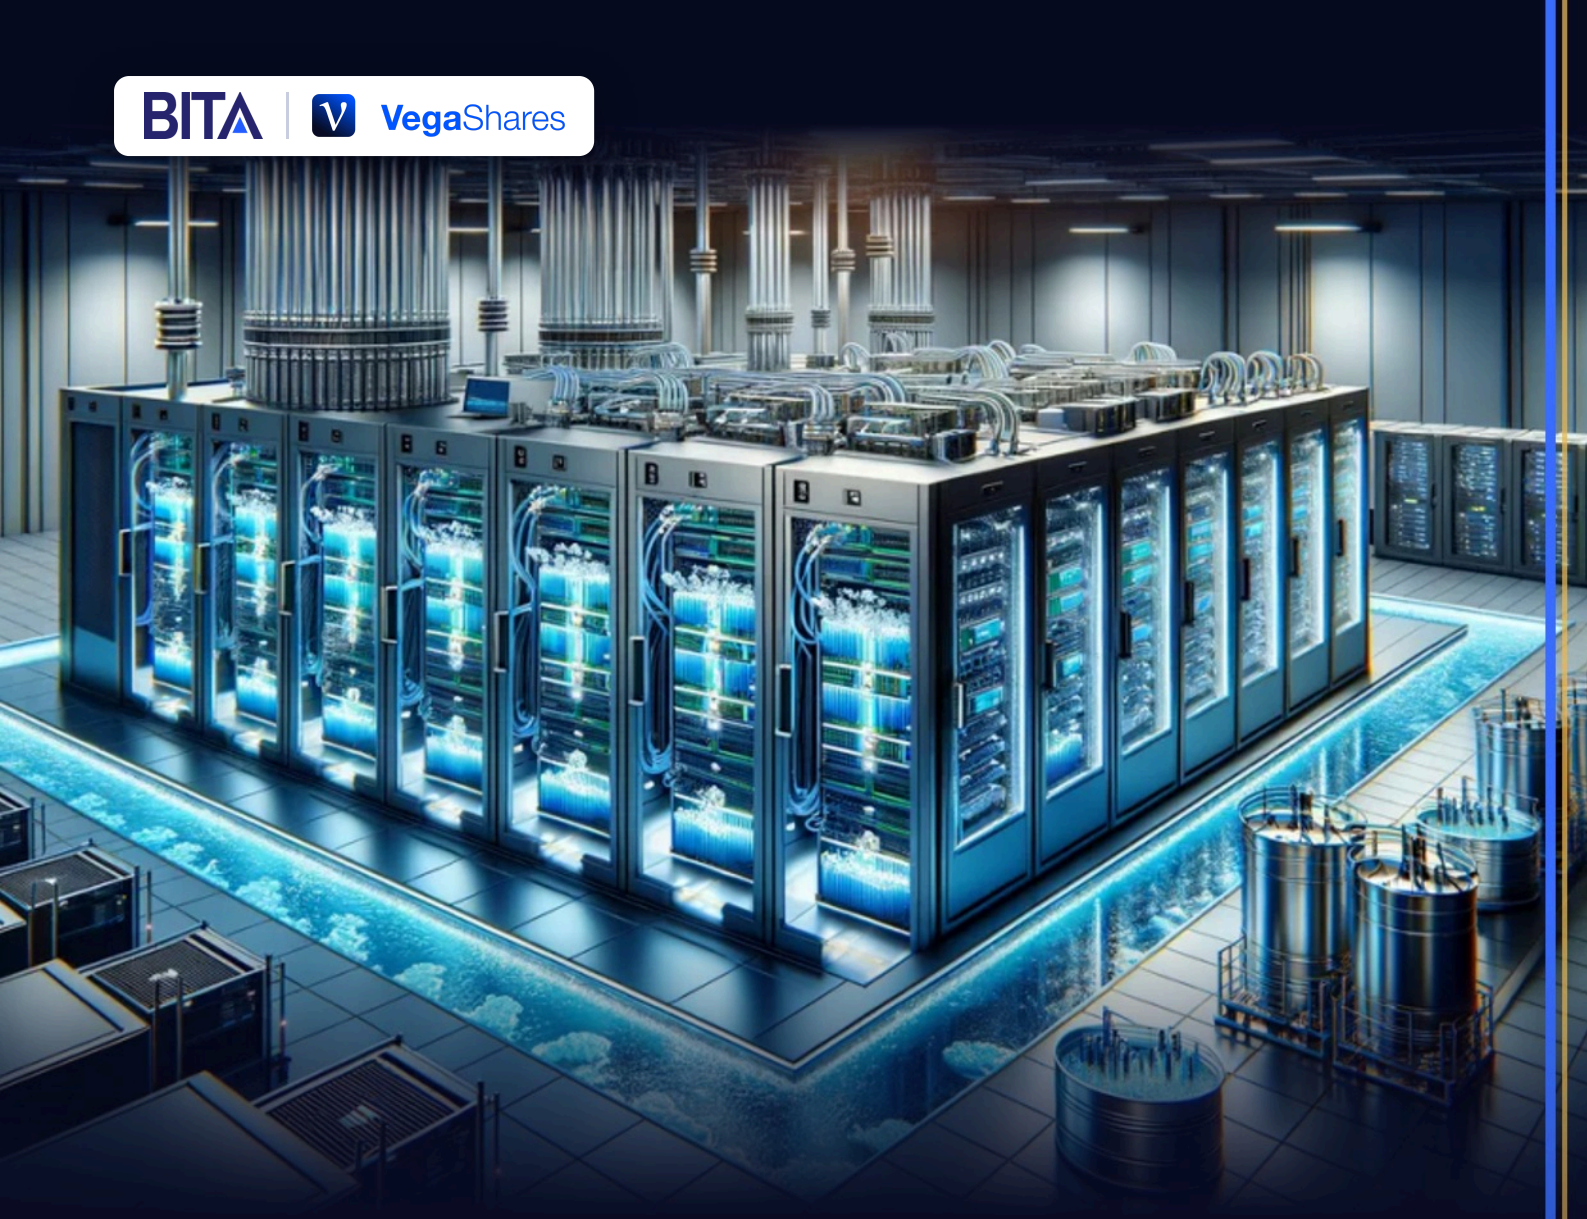
BITA | VegaShares



RESEARCH WHITE PAPER · OCTOBER 2026

Power & Heat: The Physical Limits of the AI Buildout

"Parts of the AI buildout that sit underneath the applications getting the headlines."

— Adam Stempel, Co-Founder, VegaShares ETFs

A joint research publication by BITA and VegaShares ETFs

Nasdaq: COOL · Informed by BITA

01 · EXECUTIVE SUMMARY

Power and Heat: A Distinct Way to Participate in AI Infrastructure

The AI buildout has been priced for three years through its most visible layer: compute. Graphics processing units (GPUs), high-bandwidth memory (HBM), and networking silicon have absorbed the market's attention and its capital. Underneath that layer sits a physical layer that no amount of chip supply can address on its own: every AI rack must be fed large amounts of electricity, and every watt that goes in must come back out as heat. In many new deployments, power availability and cooling capacity have become meaningful constraints on how fast new AI capacity can come online, alongside the chip supply constraints that have dominated the headlines.

The VegaShares AI Thermal, Cooling & Power Management ETF (Nasdaq: COOL) is an actively managed fund that invests in companies delivering electricity, regulating power, and removing heat from high-density AI data centers. In selecting investments, COOL's adviser considers the companies in the BITA AI Thermal, Cooling & Power Management Global Index (BAITCP), a rules-based benchmark built around the same layer, along with other companies and exposures it believes are connected to AI power and thermal management.

At a Glance

120–150 kW
Power draw of a modern Nvidia Blackwell-class AI rack, vs. 5–15 kW for a traditional enterprise rack[1]

~2,060 GW
U.S. generation and storage capacity waiting in grid interconnection queues at year-end 2025[2]

13%
Share of 2000–2020 U.S. interconnection requests operating by year-end 2025[2]

549 GW
U.S. capacity with interconnection agreements but not yet operating, year-end 2025[2]

What This Paper Covers

- How power reaches an AI chip and how heat leaves it, including the server coolant loop and the separate facility heat-rejection loop.
- How BAITCP is constructed and what its constituents' combined exposure represents, including the integrated manufacturers that make up a substantial share of the index.
- Why power and cooling have become increasingly important constraints in many AI deployments, and why that capacity is slow to add.
- A balanced view of the risks: valuation, competition, customer concentration, and project delays.

The Signal Is Already in the Earnings

Vertiv Holdings
$15.0B
Backlog at year-end 2025, up 109% YoY, with a Q4 2025 book-to-bill ratio of 2.9x.[3]

Asia Vital Components
+66%
Q2 2026 revenue growth, YoY, with gross margin rising to 32.6% from 29.8% in Q1.[4]

Schneider Electric
14–19%
Upgraded FY2026 organic adjusted EBITA growth guidance (from 10–15%).[5]

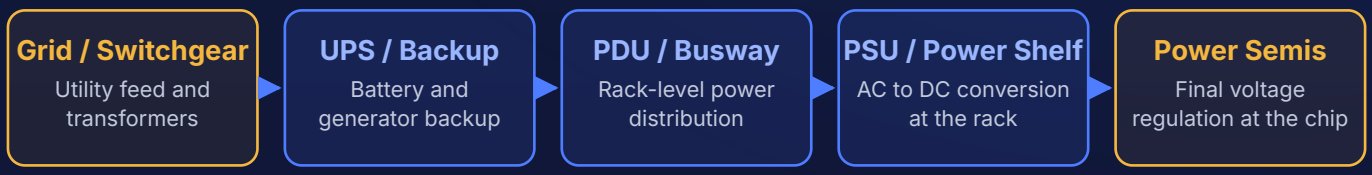
02 · FUNDAMENTALS

What Actually Powers and Cools an AI Rack?

Every AI server performs two physical jobs simultaneously: it converts incoming electricity into the precise, stable voltages a GPU needs, and it removes the heat that same GPU generates. Neither job is optional, and neither happens by accident. Each is its own chain of specialized hardware, and critically, the cooling side is really two loops: one inside the server and rack, and a separate one at the facility level.

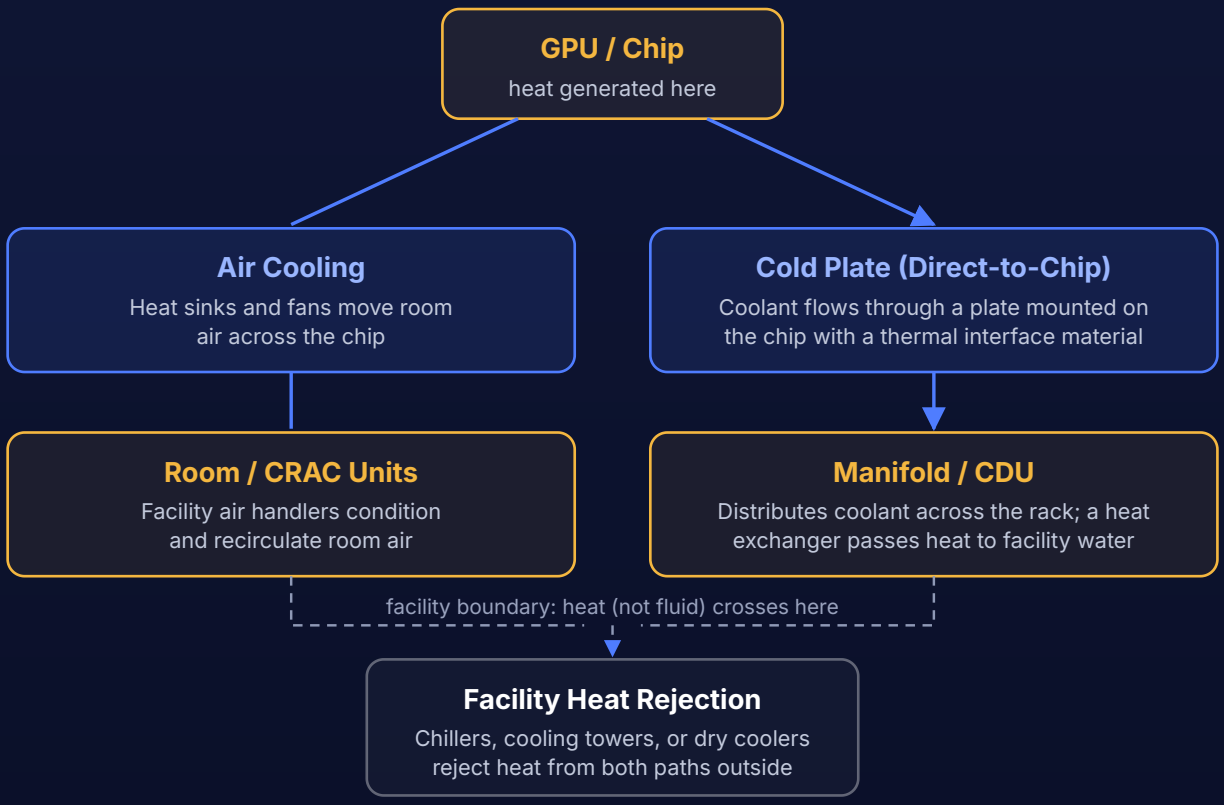
The Power Chain

Power steps down in voltage at each stage, from the utility feed to the less-than-one-volt supply a GPU core runs on.



The Cooling Side: Two Separate Loops

Air cooling and liquid cooling are not sequential stages where one feeds into the other. They are two independent paths that coexist on the same rack: a chip is cooled by moving air, by a cold plate, or often both, with each component assigned to whichever path its heat load requires. And within liquid cooling itself, the coolant that touches the server never mixes with the facility's water: the two loops exchange heat, not fluid, at the coolant distribution unit (CDU).



02 · FUNDAMENTALS (CONT'D)

Why AI Racks Need So Much More of Both

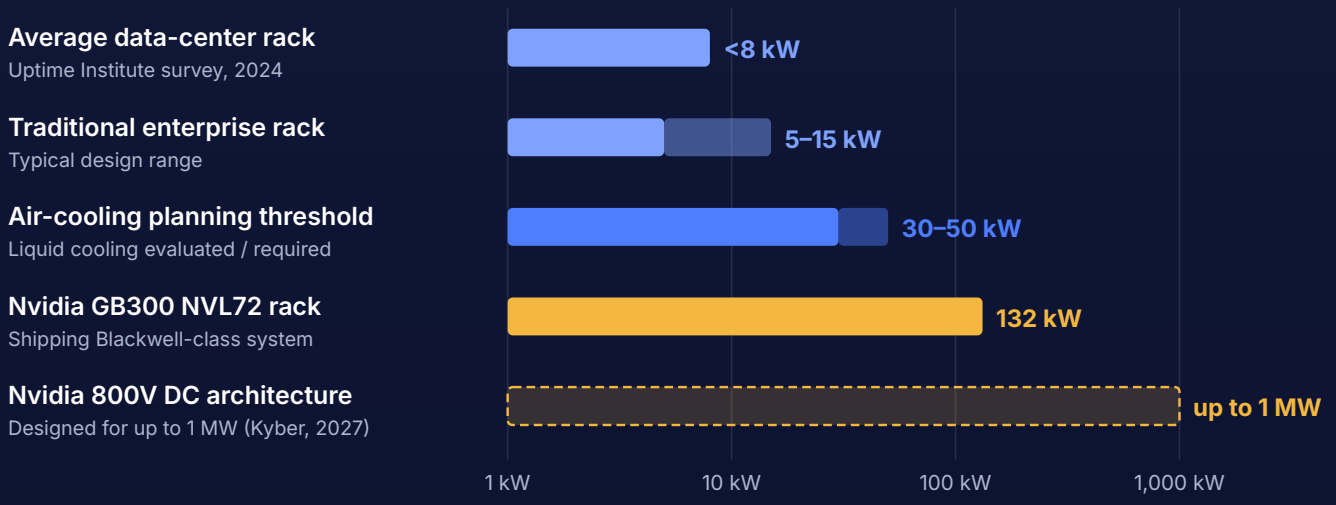
A traditional enterprise server rack draws 5–15 kilowatts and sheds that heat into moving air without much drama. A rack built around Nvidia's Blackwell-generation GPUs draws 120–150 kilowatts, an order of magnitude more, because AI training and inference pack far more silicon, at far higher power density, into the same footprint.[1] Air cannot carry that much heat away efficiently at this density, which is why liquid cooling has moved from a niche option toward a default design choice in new, high-density AI capacity.



Illustrative rendering of a liquid-cooled AI server hall.

The Power Density Ladder

Rack power by design point, on a logarithmic scale: each gridline is ten times the one before. Where a range applies, the lighter segment runs from the low end to the high end.



Sources: average rack density from Uptime Institute's 2024 Global Data Center Survey;[7] air-cooling thresholds are industry planning ranges, not hard limits;[6] GB300 NVL72 figure is HPE's nominal rack thermal design power;[1] 1 MW is the rack size Nvidia's 800V DC architecture is designed to support, with full-scale production planned alongside its Kyber rack systems in 2027, a forward-looking target rather than a shipping product.[8]

More Silicon per Rack

AI training clusters pack dozens of high-power GPUs into a single rack to minimize the latency between them, concentrating power draw and heat generation in one place rather than spreading it out.

Air Cooling's Practical Limit

Many operators begin evaluating liquid cooling as racks approach 30 kW, and at 40–50 kW cabinet heat starts to challenge the practical limits of air.[6] Beyond that, direct liquid contact with the chip becomes the more efficient path.

Power Chain Under the Same Strain

Every stage of the power chain, including PDUs, power shelves, and power semiconductors, must be re-engineered for the same density jump, not just upsized.

03 · MARKET CONTEXT

The AI Infrastructure ETF Landscape: An Overview

The thematic ETF market has moved quickly to price the AI buildout, with coverage concentrated mostly on the compute layer. GPU designers, foundries, memory makers, and networking-silicon suppliers dominate nearly every "AI infrastructure" basket on the market today. That focus captures the most visible and best-understood part of the value chain. It also leaves a gap: the physical layer of power delivery and thermal management that, increasingly, helps determine how fast new compute can actually be deployed.



A hyperscale data center campus. Power and cooling infrastructure is typically built alongside, not after, the server halls.

A Distinct Way to Participate in AI Spending

- **Concentration in chip design names:** the highest-profile "AI infrastructure" and semiconductor ETFs are typically built around the same visible compute names, rather than the power and thermal hardware that supports them.
- **Power and cooling are a different exposure, not a hidden one:** these businesses report their results publicly and are well covered by sell-side analysts. What is less common is a single-ticker product organized specifically around this layer of AI spending, distinct from chip or hyperscaler exposure.
- **A precedent from adjacent build-outs:** the 2010s cloud data-center expansion drove durable, multi-year demand for power and cooling suppliers. AI is repeating a similar pattern, at higher power density and on a faster timeline.
- **Rarely its own line item:** power delivery and thermal management are usually folded into broader "data center" or "industrials" coverage rather than analyzed as a distinct AI-infrastructure theme in its own right.

485 TWh

Estimated global data-center electricity consumption in 2025, per the International Energy Agency (IEA), which projects demand will nearly double by 2030[9]

40% by 2027

Share of existing AI data centers Gartner projects will be operationally constrained by power availability by 2027, per its November 2024 forecast[10]

04 · THE INVESTMENT CASE

The Investment Case for Power and Cooling Suppliers

A GPU shortage is, at its core, a manufacturing problem that more fab capacity can eventually solve. A grid interconnection queue is a different kind of problem: utilities and transmission operators cannot simply accelerate years of planning, permitting, and construction on demand. Power and cooling are increasingly important constraints in many AI deployments, and the companies that deliver power and remove heat from AI data centers are seeing order books and backlogs that reflect that reality.



The Grid Moves on Its Own Timeline

Approximately 2,060 gigawatts of generation and storage capacity sat in U.S. interconnection queues at year-end 2025, per Lawrence Berkeley National Laboratory. Most of it never gets built: of the capacity that requested interconnection from 2000 to 2020, only 13% had reached commercial operation by the end of 2025, while 75% had been withdrawn.[2]

A Fleet Built for a Different Era

Uptime Institute's 2024 global survey put the average data-center rack below 8 kW.[7] Most of today's installed base was designed around those densities, not the 50 kW-plus that AI workloads require, rising to 120–150 kW for Blackwell-class racks.[1]

13%

Share of U.S. capacity requesting grid interconnection in 2000–2020 that was operating by year-end 2025[2]

2.9x

Vertiv's Q4 2025 book-to-bill ratio: nearly $3 of new orders per $1 of revenue recognized[3]

549 GW

U.S. capacity with a draft or signed interconnection agreement but not yet operating, year-end 2025[2]

A Real Exposure, With Real Risks

Long-term infrastructure demand for power and cooling remains exposed to capital-spending cycles: hyperscaler capex plans can be revised, projects can be delayed or cancelled, and today's backlog growth will not repeat indefinitely. The next two pages examine the structural reasons demand has been running ahead of capacity; the Final Remarks section addresses valuation, competition, and concentration risk directly.

04 · THE INVESTMENT CASE (DEEP DIVE)

Inside the Constraint: Power

AI has turned data centers into some of the largest, most concentrated electrical loads utilities have had to connect, and getting them powered has become a genuine chokepoint in many regions.

1. The Interconnection Queue Is a Structural Constraint

A data center building can often go up faster than the grid connection it needs: the transmission and substation upgrades a large new load requires routinely outlast the building itself. The supply side is congested too: Lawrence Berkeley National Laboratory found roughly 2,060 GW of generation and storage capacity still seeking transmission interconnection at year-end 2025, and only a small fraction of capacity that enters the queue is ever built.[2]

2. Operators Are Responding by Diversifying Power Sources

Behind-the-Meter Generation

With grid capacity slow to arrive, suppliers of on-site power systems are competing for a growing share of data-center demand, and operators increasingly treat on-site generation as part of their long-term power strategy rather than relying solely on utility-delivered power.[11]

Rack Voltage Migration

Today's AI racks distribute power at 54V DC, which becomes impractical as racks pass 200 kW. Nvidia has set out an 800V DC architecture designed for racks of up to 1 MW, with full-scale production timed to its Kyber rack systems in 2027. Power semiconductors and conversion hardware must be re-engineered at every stage, not simply scaled up.[8]

3. Demand Forecasts Point to Continued Growth

75.8 GW

Projected 2026 U.S. data-center demand for grid power, per S&P Global / 451 Research[11]

~134 GW

The same outlook's projection for 2030; estimates vary meaningfully across providers[11]

up to +267%

Rise in single-month wholesale power prices versus five years earlier in some areas near data-center activity; not a national average[12]

This is the setup behind Schneider Electric's and Vertiv's results on page 12. The demand backdrop is real, but as with any infrastructure capex cycle, order books can also slow as capacity catches up or spending plans shift.

04 · THE INVESTMENT CASE (DEEP DIVE)

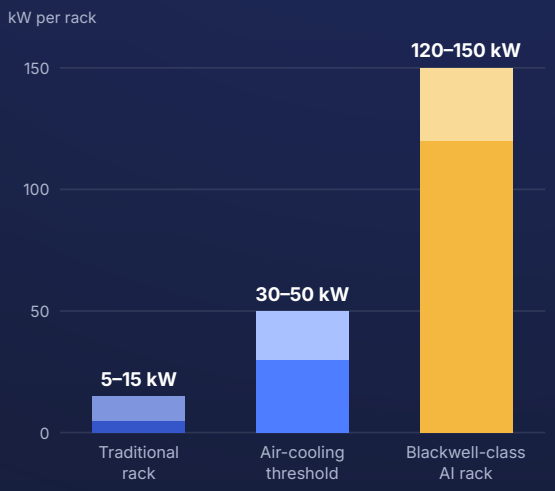
Inside the Constraint: Heat

Air cooling has a practical limit. Many operators begin evaluating liquid cooling as racks approach 30 kW, and at 40–50 kW cabinet heat starts to challenge the practical limits of air.[6] A modern AI rack running Blackwell-class GPUs draws 120–150 kW, well past that range,[1] which is why direct-to-chip liquid cooling has moved from optional to structural in new, high-density AI deployments.

Retrofitting is harder than building new. A facility engineered around 5–15 kW racks cannot simply have liquid cooling added on; it typically needs new coolant distribution units, piping, leak detection, and floor-loading capacity: capital projects measured in quarters, not weeks, layered on top of whatever power constraints already exist.

The supply chain is scaling from a small base. Cold plates, manifolds, and coolant distribution units were a niche product line three years ago. Asia Vital Components, which Nvidia named among four cold-plate suppliers for its Vera Rubin platform at the 2026 GPU Technology Conference (GTC), has told investors it is significantly expanding cold-plate and manifold output; this is management's own characterization rather than an independently audited figure.[4]

Air vs. Liquid: Rack Power Ranges



Bars are drawn to scale. The solid portion reaches the low end of each range; the lighter cap extends to the high end.[1,6]

Liquid Cooling's Adoption Curve

~30%
Share of Asia Vital Components' revenue from liquid cooling in Q1 2026, company-reported[4]

>50%
Liquid-cooling penetration in data centers that AVC management projects for 2027; an internal estimate, not independently verified[4]

1 of 4
Cold-plate suppliers Nvidia named for its Vera Rubin platform at GTC 2026, AVC among them, per press reports[4]

05 · THE INDEX

The BITA AI Thermal, Cooling & Power Management Global Index

COOL is actively managed. In selecting investments, its adviser considers the companies in the BITA AI Thermal, Cooling & Power Management Global Index (BAITCP), a rules-based benchmark built around the same power-and-thermal layer, along with other companies and exposures it believes are connected to the theme. BAITCP itself is the product of BITA's research process for mapping a fast-moving physical layer of the AI buildout into an investable universe.

Four Sub-Themes, One Physical Layer

Rather than screening for a single business description, BITA's research maps eligible companies into four sub-themes that together cover the full power-and-cooling stack:

A. Power Semiconductors & Fab Equipment

Silicon, silicon carbide, and gallium nitride power discretes, management ICs, epitaxy, substrates, and the fabrication and test equipment behind them.

B. Power Conversion & Distribution

AI server power supply units, voltage regulator modules, solid-state transformers, HVDC converters, rack-level power shelves, and uninterruptible power supplies.

C. Advanced Thermal Materials & Liquid Cooling

Thermal interface materials, heat sinks and vapor chambers, cold plates, liquid cooling pumps, manifolds, coolant distribution units, and immersion cooling tanks.

D. Integrated Systems, Software & Monitoring

Fully integrated liquid-cooled server and rack platforms, prefabricated modular power/cooling containers, and data-center infrastructure management software.

How the Universe Is Defined

Eligibility runs on a company's **BITA Thematic Exposure Score**, the share of revenue tied to the sub-themes above, built from regulatory filings, earnings transcripts, and credible news sources rather than a single reported segment line. A company qualifies if it derives at least 40% of revenue from the theme, or at least 25% while ranking among the top five companies in the universe by nominal thematic revenue, a path built specifically so a large, diversified company with a genuinely significant but non-majority AI power/cooling business can still qualify.



A live data center hall, the end destination for every constituent's products.

05 · THE INDEX

The Power & Cooling Value Chain

BAITCP's constituent count is variable by design, between 8 and 15 companies, reviewed every quarter, and currently spans five listing geographies and roughly $440 billion in aggregate float-adjusted market value.



Taiwan	46.98%	Asia Vital Components, Quanta, Wiwynn, Inventec, Auras, Nidec Chaun-Choung
USA	20.51%	Vertiv, Modine, Vicor, Navitas
France	16.74%	Schneider Electric
Canada	12.24%	Celestica
China	3.53%	Shenzhen Envicool

What the Index's Weights Actually Represent

BAITCP's narrative emphasis on thermal specialists can understate a substantial part of the portfolio: **Quanta Computer, Wiwynn, Celestica, and Inventec**, the integrated server and rack manufacturers in sub-theme D, together represent **41.22%** of index weight, the largest of the four sub-themes, and more than sub-themes A and C combined (28.78%, excluding Vertiv, which spans sub-themes B and C). These companies do not just assemble servers; they integrate power delivery and cooling hardware, sourced from across sub-themes A through C, into the finished rack platforms hyperscalers deploy. A cold plate or power shelf only reaches a data center once an integrator has designed it into a qualified rack. The following pages profile Celestica, one of the four integrated-systems constituents.

Constituents (as of September 1, 2026)

NAME	TICKER	LISTING	SUB-THEME	WEIGHT
Schneider Electric SE	SU	France	B	16.74%
Asia Vital Components	3017	Taiwan	C	16.43%
Quanta Computer Inc	2382	Taiwan	D	13.60%
Vertiv Holdings Co	VRT	USA	B/C	13.26%
Wiwynn Corp	6669	Taiwan	D	12.97%
Celestica Inc	CLS	Canada	D	12.24%
Modine Manufacturing Co	MOD	USA	C	4.24%
Shenzhen Envicool Technology-A	002837	China	C	3.53%
Inventec Corp	2356	Taiwan	D	2.41%
Vicor Corp	VICR	USA	A	1.96%
Auras Technology Co Ltd	3324	Taiwan	C	1.49%
Navitas Semiconductor Corp	NVTS	USA	A	1.05%
Nidec Chaun-Choung Technology	6230	Taiwan	C	0.08%
Total				100.00%

The 15% issuer cap is applied at each quarterly rebalancing. Between rebalancings, weights move with prices and can exceed 15%; as of September 1, 2026, Schneider Electric and Asia Vital Components were above 15% due to price movements since the last rebalance.

05 · THE INDEX

Constituent Weights & Concentration

A closer look at how BAITCP's roughly $440 billion in aggregate float-adjusted market value is distributed across its thirteen constituents, and what that concentration means for investors.

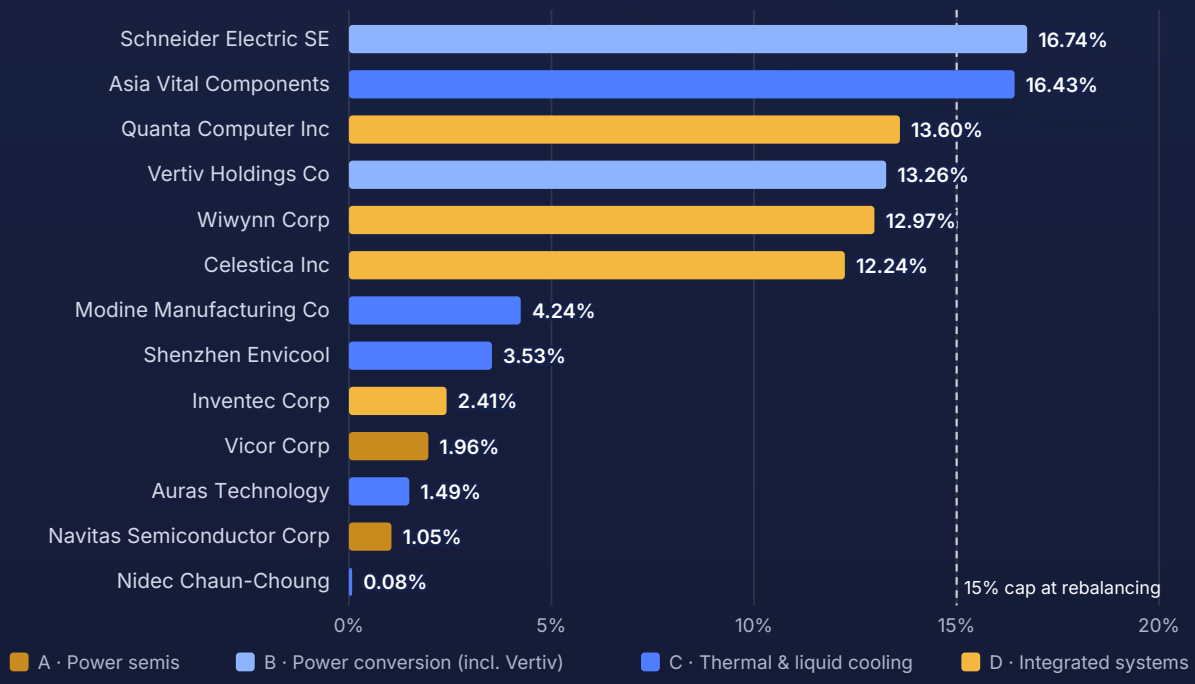
85.24%
Top 6: Schneider Electric, AVC, Quanta, Vertiv, Wiwynn, Celestica

41.22%
Integrated systems (D): Quanta, Wiwynn, Celestica, Inventec

46.98%
Taiwan: six of thirteen constituents

3.01%
Power semis (A): Vicor and Navitas Semiconductor

Weight by Constituent, September 1, 2026



That concentration reflects a supply chain in which a handful of large, publicly traded suppliers and integrators account for most of the capacity in this theme. The 15% issuer cap, applied at each quarterly rebalancing, limits single-name concentration; as a result, the largest constituents can carry similar weights even where their market values differ substantially.

For investors, this is a meaningful consideration rather than a footnote: a portfolio this concentrated, with nearly half its weight in a single geography and most of its weight in six names, carries single-name, single-sector, and single-country risk that a more diversified technology allocation would not. See Disclosures for the Fund's non-diversification and concentration risk factors.

06 · WHO IS WHO IN AI POWER & COOLING

Global Power & Thermal Leaders

Two of BAITCP's largest constituents sit at the center of the power-and-cooling buildout because they sell into both sides of it: power delivery and thermal management, increasingly packaged as a single integrated system. Figures below are as of each company's most recent reporting as of September 2026.

Schneider Electric (France)

BAITCP weight 16.74%

14–19%
Upgraded FY2026 organic adjusted EBITA growth guidance

+22%
H1 2026 organic adjusted EBITA growth, YoY

Schneider raised its full-year 2026 organic adjusted EBITA growth guidance to 14–19%, from an earlier 10–15%, after H1 2026 revenue reached a record €21.2 billion (+14% organic) and adjusted EBITA reached €4.1 billion (+22% organic), with margin expanding 120bps* to 19.3%. Its Motivair subsidiary supplies liquid cooling, and its core Energy Management segment supplies rack power distribution and uninterruptible power supplies (UPS); management increasingly frames the company as an end-to-end data-center infrastructure provider.[5]

Vertiv Holdings (United States)

BAITCP weight 13.26%

$15.0B
Backlog at year-end 2025, up 109% YoY

$14.0B
FY2026 net sales guidance midpoint, raised in July 2026

Vertiv's Q4 2025 organic orders rose approximately 252% year-over-year, lifting backlog to $15.0 billion (+109% YoY) with a 2.9x book-to-bill ratio for the quarter. Year-end 2025 is Vertiv's most recent backlog disclosure; in February 2026 the company stopped reporting orders and backlog with quarterly earnings. Full-year 2025 adjusted free cash flow reached $1.9 billion, and in July 2026 management raised full-year 2026 net sales guidance to $13.8–$14.2 billion.[3]

Why Integration Is Winning

Both companies are converging on the same pitch: a hyperscaler building a new AI hall does not want to separately source power distribution, backup, and cooling from different vendors and integrate them on-site. Schneider and Vertiv are positioning to sell power and thermal as one designed system, which can raise switching costs and extend order visibility, though it also means both companies' results are exposed to the same AI-capex cycle as the hyperscalers they serve.

*bps = basis points, one hundredth of one percent (0.01%)

06 · WHO IS WHO IN AI POWER & COOLING

Integrated Systems: Where Power and Cooling Get Built In

Sub-theme D, integrated server and rack manufacturers, is BAITCP's single largest sub-theme by combined weight. These companies do not make cold plates or power semiconductors themselves; they design and build the rack-level platforms that integrate components from across the other three sub-themes into what a hyperscaler actually deploys.

celestica

Celestica (Canada)

BAITCP weight 12.24%

+62%
Q2 2026 revenue growth, YoY

$20.5B
Raised FY2026 revenue outlook (+65% YoY)

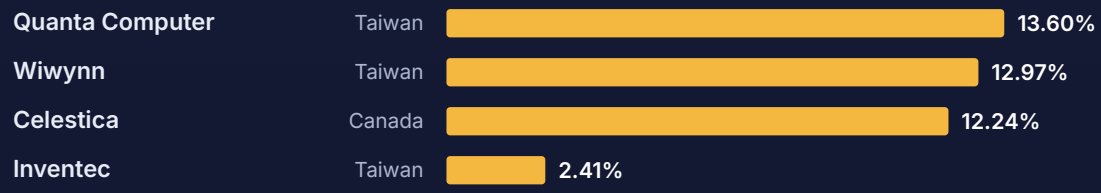
+58%
Q2 2026 Hardware Platform Solutions revenue growth, YoY

Celestica's Q2 2026 revenue reached $4.70 billion, up 62% year-over-year, with its Connectivity & Cloud Solutions segment up 84%. Its Hardware Platform Solutions business, which designs and manufactures hardware platforms for hyperscale customers, generated approximately $1.9 billion in the quarter, up 58% year-over-year; Celestica attributed that growth primarily to networking switch programs. Management raised its full-year 2026 outlook twice during the year, to $20.5 billion in revenue and $11.30 in adjusted EPS. Three customers accounted for 32%, 17%, and 14% of Q2 2026 revenue.[13]

What an Integrator Actually Does

A cold plate, a power shelf, and a thermal interface material are each engineered by a specialist supplier. None of them reaches a data center on its own. Quanta, Wiwynn, Celestica, and Inventec are the companies that take components sourced across sub-themes A through C and turn them into a qualified, deployable rack: they design the mechanical and electrical integration, validate thermal and power performance together, and manufacture at the scale hyperscalers require. That integration role is also why these companies can be more exposed to customer concentration risk than a pure-play component supplier: their revenue depends heavily on a small number of hyperscaler design wins and program ramps, which can also be delayed or re-scoped.

Sub-Theme D Constituents: 41.22% of BAITCP



Index weights as of September 1, 2026.

06 · WHO IS WHO IN AI POWER & COOLING

Thermal Management & Power Semiconductors

Asia Vital Components (Taiwan)

BAITCP weight 16.43%

+66%
Q2 2026 revenue growth, YoY

32.6%
Q2 2026 gross margin, up from 29.8% in Q1

AVC's Q2 2026 revenue reached approximately NT$49.1 billion, up 66% year-over-year, and first-half revenue rose 85%. Liquid cooling was about 30% of total sales in Q1 2026. Management has described 2026 as the start of a "liquid cooling era" and has stated it is expanding production capacity significantly; these are company characterizations rather than independently audited figures.[4]

Vicor Corporation (United States)

BAITCP weight 1.96%

+26.9%
Q2 2026 revenue growth, sequential

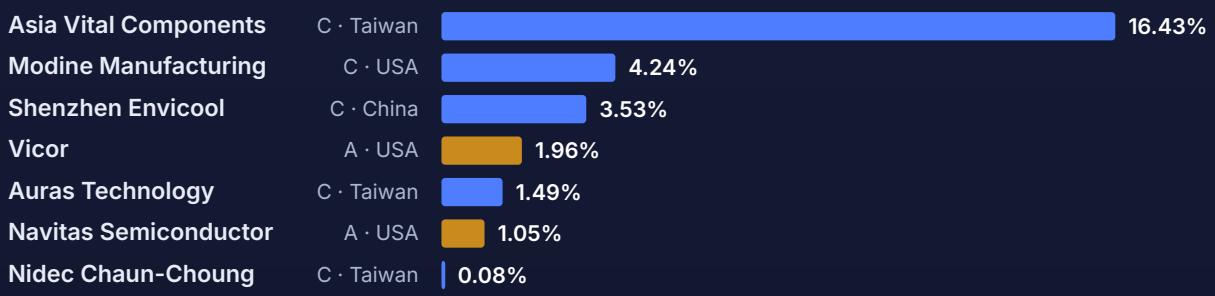
$600M+
Expected FY2026 revenue

Vicor's Q2 2026 revenue reached $143.4 million, up 26.9% sequentially, with EPS of $1.04; the quarter included $15 million of royalty income under a new license agreement, which management expects to contribute $5 million in Q3. Vicor's second-generation Vertical Power Delivery technology targets power regulation directly at the processor, and management said the company is approaching capacity utilization and will need a second fab to reach higher volumes.[14]

Two Different Geographies, One Supply Chain

AVC is based in Taiwan, part of the geographic concentration that defines sub-theme C. Vicor is a U.S. company; together they illustrate that thermal and power-semiconductor exposure in this index is not confined to a single country, even though the integrated-systems manufacturers in sub-theme D are heavily Taiwan-based. BAITCP's remaining constituents fill out these sub-themes further: **Modine Manufacturing** and **Shenzhen Envicool** supply liquid-cooling systems alongside AVC and Auras Technology, and **Navitas Semiconductor**, which in August 2026 agreed to acquire power-delivery specialist Claros in a transaction valued at up to approximately $232.8 million, including milestone-based consideration (the deal is expected to close by year-end 2026), rounds out the power-semiconductor exposure alongside Vicor.[15]

Thermal (C) and Power-Semiconductor (A) Constituents



Index weights as of September 1, 2026. Sub-theme C totals 25.77% excluding Vertiv (B/C); sub-theme A totals 3.01%.

07 · FINAL REMARKS

Final Remarks

For three years, the market has priced the AI buildout largely through its most visible layer: compute. Every AI rack also depends on two less visible chains: one that delivers electricity in the voltages and volumes a GPU needs, and one that removes the heat that same GPU generates. Both chains are under real strain in many deployments, evidenced by congested grid interconnection queues, a data-center fleet built for lower rack densities than AI workloads now require, and strong order growth at the companies supplying this layer.

COOL gives investors active, targeted access to this layer through a single ticker, informed by the research embedded in BAITCP. The strongest version of this argument is straightforward: power delivery and thermal management offer a distinct way to participate in AI infrastructure spending, separate from chip or hyperscaler exposure. It is a different claim from saying the market has overlooked these businesses, which is harder to substantiate given how closely sell-side analysts already cover Schneider Electric, Vertiv, and similar names.

Risks Worth Weighing

Valuation and Cyclicality

Several constituents have re-rated substantially on AI-driven order growth. Long-term infrastructure demand remains exposed to capital-spending cycles, and a slowdown in hyperscaler capex, a pause in AI data-center construction, or a normalization of order growth could compress both earnings growth and multiples at the same time.

Competition and Margin Pressure

Strong demand is attracting new entrants and capacity expansion across cooling and power hardware. Increased competition, particularly from lower-cost manufacturers entering liquid cooling, could pressure pricing and margins even if unit demand keeps growing.

Customer Concentration

A small number of hyperscalers drive a large share of incremental demand across this theme. The integrated-systems manufacturers profiled on page 13 are especially exposed to program-level decisions by a handful of customers.

Order Cancellations and Project Delays

Backlogs and book-to-bill ratios reflect orders placed, not revenue guaranteed to be recognized on schedule. Data-center construction timelines, permitting delays, and project cancellations have historically caused reported backlogs to convert unevenly into revenue.

Strong industry demand does not automatically translate into attractive shareholder returns; execution, pricing discipline, and capital allocation still matter. For investors who have already priced the chips, COOL is one way to gain exposure to what keeps them powered and cool, evaluated alongside the risks above rather than in place of them.

08 · DISCLOSURES & SOURCES

Disclosures & Sources

Investing involves risk, including possible loss of principal. Fund objectives, risks, charges, and expenses should be considered carefully before investing. For a prospectus or summary prospectus with this and other important information about the Fund, call (888) 862-3299 or visit VegaSharesETFs.com.

COOL is an actively managed ETF. In selecting investments, the adviser considers companies included in the BITA AI Thermal, Cooling & Power Management Global Index (BAITCP), along with other companies and investment exposures it believes are connected to AI power and thermal management; the Fund is not required to invest only in index constituents, and its holdings can differ materially from the Index at any time. BAITCP is a rules-based index owned, calculated, administered, and disseminated by BITA GmbH, which is not affiliated with VegaShares ETFs or the Fund's adviser. The Fund is non-diversified and concentrated in industries related to AI power and thermal management, which may make it more volatile than a broadly diversified fund.

Statements regarding future industry growth, market share, or production capacity, including those attributed to company management, reflect third-party estimates or management's own characterization of its business and are not independently verified; they are not guarantees of future results. Company figures are as reported by each company; figures attributed to press reports or earnings summaries are identified as such below.

Numbered Sources

1 Hewlett Packard Enterprise, “NVIDIA GB300 NVL72 by HPE” product documentation (rack thermal design power 132 kW nominal, approximately 155 kW peak), 2026.

2 Lawrence Berkeley National Laboratory, “Queued Up: 2026 Edition,” June 2026, data through December 31, 2025 (~2,061 GW active; 549 GW with interconnection agreements not yet operating; 13% of 2000–2020 requested capacity operating by end-2025). The dataset covers generation and storage requests only, not load interconnection.

3 Vertiv Holdings Co, Fourth Quarter and Full Year 2025 Results, February 11, 2026 (backlog $15.0B, +109% YoY; Q4 2025 book-to-bill ~2.9x; FY2025 adjusted free cash flow $1,887M), and earnings call, February 11, 2026 (discontinuation of quarterly orders and backlog reporting); Second Quarter 2026 Results, July 29, 2026 (FY2026 net sales guidance $13.8–$14.2B).

4 Asia Vital Components, Q1 2026 results (revenue NT$49.04B, +110% YoY; gross margin 29.77%; liquid cooling ~30% of sales) and Q2 2026 results and earnings call, August 2026, as summarized by Quartr and Investing.com (revenue ~NT$49.1B, +66% YoY; gross margin 32.57%; H1 revenue +85%); DIGITIMES, March 2026, and TVBS News, April 16, 2026 (Nvidia's GTC 2026 cold-plate supplier selection). Market-share, capacity, and penetration figures reflect management's own statements and are not independently audited.

5 Schneider Electric, Half-Year 2026 Results, July 30, 2026 (FY2026 organic adjusted EBITA growth guidance raised to 14–19% from 10–15%). EBITA is earnings before interest, taxes, and amortization.

6 Chatsworth Products, “When Do You Need Liquid Cooling for AI Infrastructure? Key Thresholds and Warning Signs,” June 9, 2026 (planning thresholds of ~30 kW and 40–50 kW per cabinet; not universal limits).

7 Uptime Institute, 2024 Global Data Center Survey (average rack density below 8 kW), as reported by Upsite Technologies.

8 NVIDIA, “NVIDIA 800 V HVDC Architecture Will Power the Next Generation of AI Factories,” NVIDIA Technical Blog, May 20, 2025 (54 VDC rack distribution today; 800 VDC designed for racks up to 1 MW; full-scale production alongside Kyber rack-scale systems in 2027).

9 International Energy Agency, Key Questions on Energy and AI, 2026 (estimated 485 TWh global data-center electricity consumption in 2025; demand projected to nearly double by 2030).

10 Gartner, “Gartner Predicts Power Shortages Will Restrict 40% of AI Data Centers By 2027,” press release, November 12, 2024.

11 S&P Global Commodity Insights, “Data center grid-power demand to rise 22% in 2025, nearly triple by 2030,” October 2025 (451 Research outlook: 75.8 GW in 2026; 134.4 GW in 2030; on-site power suppliers competing for data-center demand). Forecasts vary meaningfully across research providers.

12 Bloomberg News, “AI Data Centers Are Sending Power Bills Soaring,” September 29, 2025 (wholesale prices up to 267% higher for a single month than five years earlier in some areas near significant data-center activity).

13 Celestica Inc, Q2 2026 earnings release and Form 10-Q, July 2026; Q1 2026 earnings release, April 2026.

14 Vicor Corporation, Q2 2026 results, July 21, 2026, and Q2 2026 earnings call transcript, July 21, 2026.

15 Navitas Semiconductor, Form 8-K and press release announcing the definitive agreement to acquire Claros, August 2026 (up to ~$232.8M, of which ~$216.0M payable at closing; remainder contingent on milestones over two years).

09 · ABOUT

About VegaShares & BITA

About VegaShares

VegaShares specializes in derivatives-based ETFs and other innovative investment strategies. Developed by institutional experts, VegaShares blends quantitative research with disciplined risk management to create liquid, exchange-traded tools for modern investors seeking efficiency, precision, and performance.

About BITA

BITA is a Germany-based Fintech that provides enterprise-grade indexes, data, and infrastructure to institutions operating in the passive and quantitative investment space. Thanks to BITA's innovative index software, designed to outperform other existing solutions in terms of flexibility and speed, BITA is able to provide independent, methodologically-sound indexes that are both investable and replicable by customers and stakeholders. All of BITA's methodologies and processes are completely transparent and available publicly.

The German Federal Financial Supervisory Authority (BaFin) supervises BITA with respect to EU Paris-aligned and EU Climate Transition benchmarks; BITA voluntarily applies the wider EU Benchmark Regulation framework across its index business as a matter of good practice.

10 · FUND INFORMATION

COOL Holdings & Fund Risks

COOL Top 10 Holdings as of 10/15/2026

HOLDING	% OF NET ASSETS
First American Government Obligations Fund	21.49%
Asia Vital Components	10.17%
Vertiv Holdings Co	8.35%
Vicor Corp	6.34%
Auras Technology	6.33%

HOLDING	% OF NET ASSETS
Sunonwealth Electric Machine Industry	5.75%
Fositek Corp	5.68%
Sanyo Denki Co Ltd	5.16%
Nidec Chaun-Choung Technology	4.53%
Advanced Energy Industries Inc	4.44%

Holdings are subject to change and should not be considered a recommendation to buy or sell any security.

Fund Risks & Disclosures

Before investing, carefully consider the fund's investment objectives, risks, and charges and expenses. The prospectus and summary prospectus contain this and other important information and may be obtained by visiting VegaSharesETFs.com or calling 1-888-862-3299. Read it carefully before investing.

An investment in the fund is subject to risks, and you could lose money on your investment in the fund. There can be no assurance that the fund will achieve its investment objective. Your investment in the fund is not a deposit in a bank and is not insured or guaranteed by the Federal Deposit Insurance Corporation (FDIC) or any other government agency. The risks associated with an investment in the fund can increase during times of significant market volatility. The fund also has specific principal risks, which are described below. More detailed information regarding these risks can be found in the fund's prospectus.

Principal Risks. You could lose money on an investment in the Fund.

Datacenter Industry Risk. The datacenter industry faces numerous challenges that could significantly impact the financial performance of companies operating within this sector.

Technology Hardware, Storage & Peripherals Industry Risk. Technology hardware, storage and peripherals companies can be significantly affected by competitive pressures, aggressive pricing, technological developments, changing domestic demand, the ability to attract and retain skilled employees and availability and price of components.

Semiconductors & Semiconductor Equipment Industry Risk. Semiconductor companies may face intense competition, both domestically and internationally, and such competition may have an adverse effect on their profit margins.

Artificial Intelligence Risk. Issuers engaged in artificial intelligence typically have high research and capital expenditures and, as a result, their profitability can vary widely, if they are profitable at all.

Information Technology Sector Risk. The information technology sector includes companies engaged in internet software and services, technology hardware and storage peripherals, electronic equipment instruments and components, and semiconductors and semiconductor equipment, among other things.

Concentration and Non-Diversification Risk. Focused exposures may produce greater volatility than a more diversified portfolio.

Derivatives Risk. Derivatives are financial instruments that derive value from the underlying reference asset or assets, such as stocks, bonds, or funds (including ETFs), interest rates or indexes. The Fund's investments in derivatives may pose risks in addition to, and greater than, those associated with directly investing in securities or other ordinary investments, including risk related to the market, imperfect correlation with underlying investments or the Fund's other portfolio holdings, higher price volatility, lack of availability, counterparty risk, liquidity, valuation and legal restrictions.

High Portfolio Turnover Risk. The Fund's strategy is linked to an Index that is expected to have a high portfolio turnover rate. As a result, the Fund is likewise expected to frequently trade all or a significant portion of the securities in its portfolio.

Models and Data Risk. The composition of the Fund's portfolio is heavily dependent on investment models developed by the Sub-Adviser as well as information and data supplied by third parties ("Models and Data").

Foreign Securities Risk. Foreign investments are subject to currency, depositary receipt, and political risks.

IPO Risks. The Fund may purchase securities of companies that are offered in an IPO.

SPAC and De-SPAC Risks. The Fund may invest in securities of companies that have recently completed IPOs and may have become publicly traded through transactions involving SPACs or de-SPAC transactions.

Limited History of Operations Risk. The Fund is a new ETF and has a limited history of operations for investors to evaluate.

Foreside Fund Services, LLC is the distributor of the VegaShares ETFs. Foreside is not affiliated with VegaShares ETFs.



www.bita.io · info@bita.io
www.vegasharesetfs.com · (888) 862-3299