Energy Crunch Opportunities: Balancing AI Innovation and Data Center Demands

Calamos Antetokounmpo Global Sustainable Equities ETF (SROI), Calamos Antetokounmpo Sustainable Equities Fund (SROIX)
Jim Madden, CFA, Tony Tursich, CFA, and Beth Williamson

As an investment team we want to provide value to our shareholders foremost. We do this by constructing investment portfolios of durable companies that make business sense and, in our opinion, address the sustainable risks and opportunities of the future. While many companies are still defining their business case for AI, the possibilities appear limitless to us. However, data center advancements that mitigate energy and resource use will be critical to the viability of AI's rapid progression, and this is where we see tremendous opportunity.

All online interactions depend on a foundation of data stored in distant servers. Those servers, stacked together in data centers around the world, require a lot of energy. Currently, data centers account for about 1.0%–1.5% of the world's electricity use[1], and this figure is expected to increase to 3%–4% by the end of the decade[2], in large part due to the exploding boom in artificial intelligence (AI).

In our view, AI has emerged as a transformative force, changing how we process, analyze, and utilize data across all industries. AI is advancing so quickly that the World Economic Forum reports the computational power required to sustain AI's rise doubles roughly every 100 days[3].

And AI computations are far more energy-intensive than conventional internet queries. For example, ChatGPT queries are 6x to 10x more power hungry than traditional Google searches[4].

POWER CONSUMPTION PER QUERY/SEARCH (WH)

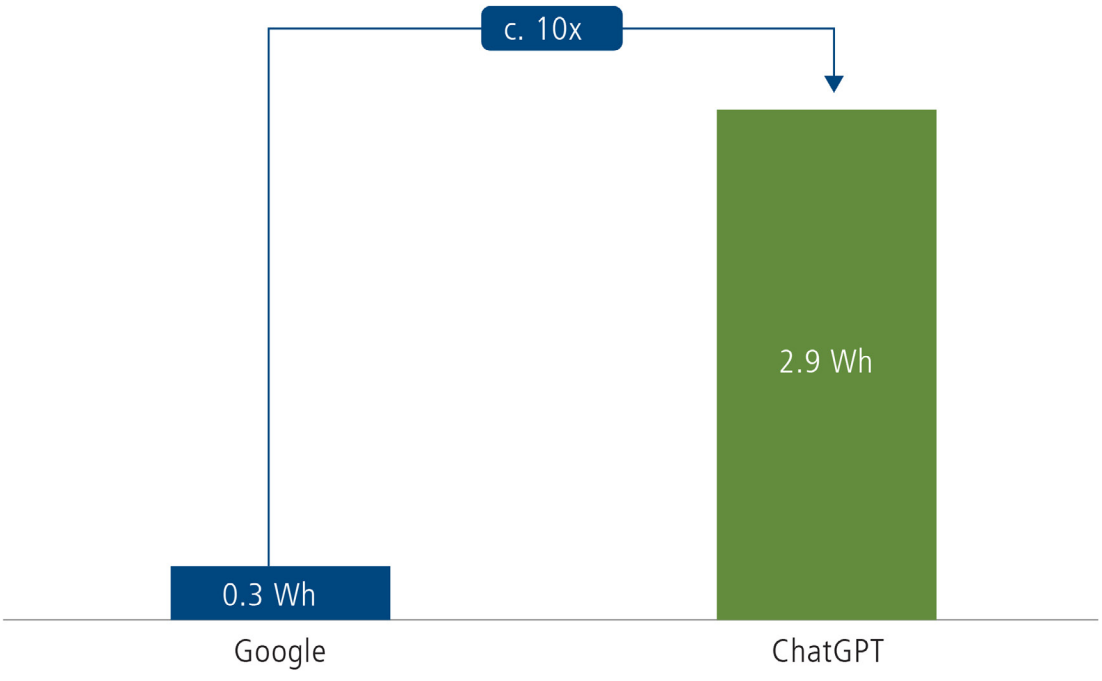

Source: Google, SemiAnalysis

According to a Schneider Electric white paper[5], AI represents about 4.5 GW of power consumption today and is projected to grow at an annual rate of 25% to 33%, resulting in a total consumption of 14 GW to 18.7 GW by 2028.

SCHNEIDER ELECTRIC ESTIMATE	2023	2028
Total data center power consumption	57 GW	93 GW
AI power consumption	4.5 GW	14.0–18.7 GW
AI power consumption (% of total)	8%	15%–20%
AI workload (Training vs Inference)	20% Training/80% Inference	15% Training/85% Inference

Source: Schneider Electric

AI WORKLOAD DEMANDS ARE DAUNTING

As referenced in the above table, the AI workload is completed in two key stages (training and inference). Both stages impact the environment mainly through energy use and water consumption. At present, the environmental footprint of AI is split, with training responsible for approximately 20% and inference 80%[6].

During the training phase, AI models learn patterns by digesting vast amounts of data, requiring significant amounts of energy. For example, the graphics processing units (GPUs) that trained GPT-3 (the precursor to ChatGPT) are estimated to have consumed 1,300 megawatt-hours of electricity, roughly equal to that used by 1,450 average U.S. households per month[7]. These models also require water for cooling (Scope 1) as well as water for power generation and manufacturing (Scope 2 and 3 accordingly).

According to JPMorgan's ChatESG[8], "Training GPT-3 in Microsoft's state-of-the-art U.S. data centers evaporate 700,000 liters of clean freshwater."

Once trained, AI models step into the inference phase to run live. While the inference phase requires less energy and water because fewer computations are involved, over time this phase is the largest contributor to emissions[9].

NEW TOOLS CAN MITIGATE AI WORKLOADS

Capping Power. To reduce energy consumption across AI workloads, manufacturers and developers alike are researching methods to limit the amount of power a GPU can draw and improve workload accuracy. "Capping power" is one technique that can be employed at data centers via software. It involves setting limits on the power consumption of hardware components such as GPUs or CPUs to manage energy usage. According to MIT's Lincoln Laboratory[10], power capping GPUs during AI model training can result in a 12%–15% reduction in energy use. The downside, capping power can increase the task time by approximately 3%. However, given a model's training duration of days, weeks or months, this time increase is negligible.

Weeding Out Underperformance. Also, during training, AI developers can focus on improving accuracy. By analyzing the rate at which the model learns, developers can stop underperforming models early. Referencing Lincoln Laboratory's research studies again, "… early stopping led to dramatic savings: an 80% reduction in the energy used for model training."

Optimizing Hardware. To improve energy efficiency during the inference stage, an optimizer can be utilized to match an AI model with the most carbon-efficient mix of hardware for example, high-power GPUs for the computationally intense parts of inference and low-power central processing units (CPUs) for the less-demanding aspects can decrease energy use by 10% to 20% without compromising performance[11].z

Market Application

Microsoft is one example of a company that has adopted tools to lessen AI workloads. According to the company, as of June 2023, Microsoft deployed its power capping system to millions of servers across the company's data centers thereby freeing up hundreds of MWs of harvested power[12]. This capping system also allowed Bing and Bing Ads to safely enhance performance by maximizing air intake, also known as a turbo boost, resulting in performance improvements of ~20%.

Microsoft has also developed custom data center chips like Azure Maia, also known as Maia 100, an AI-optimized GPU designed for running complex AI workloads. It's built on a 5nm node and optimized for scalability and sustainability, with features like dynamic power optimization and liquid cooling.

DATA CENTER INFRASTRUCTURE INNOVATIONS MAY BE VITAL TO SUCCESS

The energy needs of data centers are driven by computing (40% of electricity demand) and cooling (40%). The remaining 20% is divided between power supply, storage, and communication equipment[13]. As such, data center infrastructure improvements can also play a significant role in reducing the environmental impacts of data centers.

Since cooling accounts for 40% of a data center's energy needs, efficient cooling is a top priority. The shift from air cooling to liquid cooling is a potential infrastructure innovation many tout as transformational. In the direct-to-chip liquid cooling approach, "… a cooling fluid is circulated through the servers to absorb and dissipate heat, and is quickly gaining popularity as a more effective way to handle the concentrated heat generated by AI clusters."[14] Compared with air cooling, liquid cooling consumes 10% less energy[15], improves power utilization, and reduces water usage.

Market Application

NVIDIA is actively supporting direct-to-chip liquid cooling for its high-performance data center GPUs. They have released their first data center PCIe GPU, which utilizes this cooling method. Additionally, NVIDIA plans to continue supporting liquid cooling in their GPUs and HGX platforms. In sum, this approach enhances efficiency, sustainability, and optimal system performance for AI workloads. Furthermore, NVIDIA was awarded $5 million by the United States Department of Energy (DOE) to develop a cooling solution combining two-phase direct-to-chip and immersion cooling techniques using environmentally compliant refrigerants[16].

HYPERSCALERS MUST BE SUSTAINABILITY INNOVATORS

Approximately 8,000 data centers currently operate in the world. The US leads with a third of the total data centers, followed by Europe (16%) and China (10%). Over the last decade, the nature of these data centers has evolved, moving from "traditional" to "Cloud" and "Hyperscale"[17].

CLOUD & HYPERSCALE DATA CENTERS BECAME THE STANDARD OVER THE DECADE

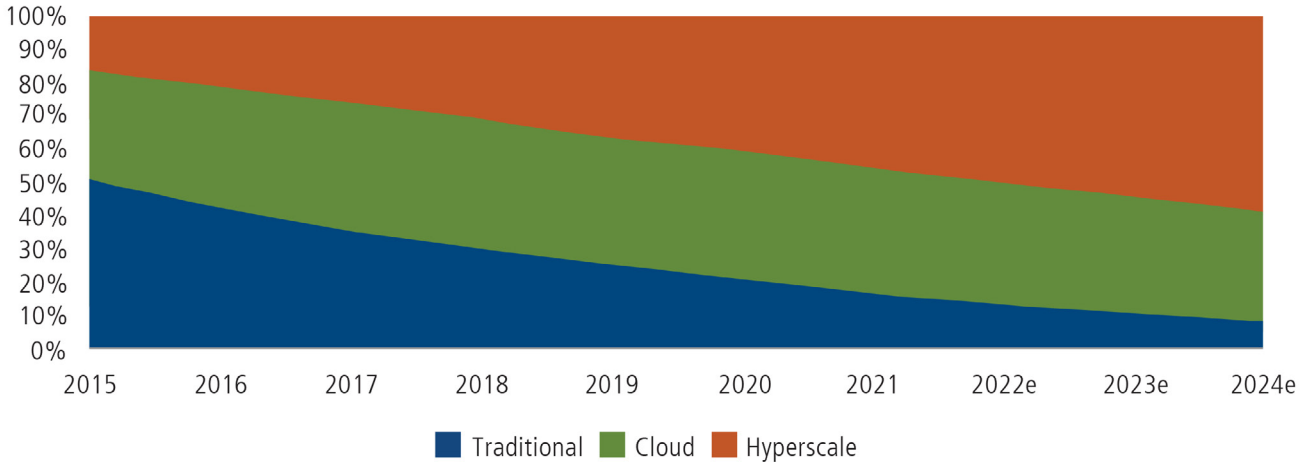

Source: JP Morgan estimates, IEA

With the prominence of hyperscalers, (i.e, the biggest data center owners—Google, Microsoft and Amazon); data centers saw an uptick in efficiency as all of the aforementioned companies have set climate goals and face internal and external pressure to deliver on them. But the rise of AI is jeopardizing these corporate goals. Current practices of sourcing renewable energy or utilizing carbon credits/offsets are no longer sufficient.

Market Application

Alphabet Inc. In addition to ensuring the use of high-efficiency hardware, resilient power, and cooling systems to improve AI workloads, hyperscalers must also ensure the procurement of renewable energy sources. To do so, many are following Google's lead to "load shifting." Rather than relying solely on the grid's mix of fossil fuels and renewables, hyperscalers are trying to shift, daily or even hourly, data center operations around the world to access excess renewable energy production operations across time zones. Google has taken a pioneering step by aligning its data center power usage with zero-carbon sources on an hourly basis. However, achieving uninterrupted clean energy remains elusive.

INFORMATION IS POWER

In our opinion, AI is a transformative technology, but its use is directly responsible for an uptick in carbon emissions and the consumption of millions of gallons of fresh water. Yet it can also be a positive, improving building efficiency, health care and climate modeling.

What is apparent is that the development of AI cannot come at the expense of our planet.

As such, many government representatives in the US and abroad are working to develop a standardized system for reporting AI impacts on society and the environment. Leading the way is the International Organization for Standardization (ISO) who will be issuing criteria for "sustainable AI," which will include standards for measuring energy efficiency, raw material use, transportation, and water consumption, as well as practices for reducing AI impacts throughout its life cycle. "The ISO wants to enable AI users to make informed decisions about their AI consumption[18]."

Calamos' Sustainable Equities team recognizes the benefits and challenges AI brings. We will continue to seek global leaders who are working to advance this long-term secular trend, sustainably.

[1] Leffer, Lauren. "The AI Boom Could Use a Shocking Amount of Electricity." Scientific American. https://www.scientificamerican.com/article/the-ai-boom-could-use-a-shocking-amount-of-electricity/

[2] "AI is poised to drive 160% increase in data center power demand." Goldman Sachs. https://www.goldmansachs.com/intelligence/pages/AI-poised-to-drive-160-increase-in-power-demand.html

[3] "How to manage AI's energy demand — today, tomorrow and in the future." World Economic Forum. https://www.weforum.org/agenda/2024/04/how-to-manage-ais-energy-demand-today-tomorrow-and-in-the-future/

[4] "AI, data centers and the coming US power demand surge." Goldman Sachs. https://www.goldmansachs.com/intelligence/pages/gs-research/generational-growth-ai-data-centers-and-the-coming-us-power-surge/report.pdf

[5] Avelar, Victor, et al. "The AI Disruption: Challenges and Guidance for Data Center Design." Energy Management Research Center. Schneider Electric. https://www.se.com/ww/en/insights/electricity-4-0/digitalization/the-ai-disruption.jsp

[6] "How to manage AI's energy demand — today, tomorrow and in the future." World Economic Forum. https://www.weforum.org/agenda/2024/04/how-to-manage-ais-energy-demand-today-tomorrow-and-in-the-future/

[7] Foy, Kylie. "AI models are devouring energy. Tools to reduce consumption are here if data centers will adopt.".MIT Lincoln Laboratory. https://www.ll.mit.edu/news/ai-models-are-devouring-energy-tools-reduce-consumption-are-here-if-data-centers-will-adopt

[8] JP Morgan, ChatESG: Why is AI so thirsty? Water use by data centers 101

[9] Foy, Kylie. "AI models are devouring energy. Tools to reduce consumption are here if data centers will adopt." MIT Lincoln Laboratory. https://www.ll.mit.edu/news/ai-models-are-devouring-energy-tools-reduce-consumption-are-here-if-data-centers-will-adopt

[10] IBID

[11] Law, Marcus. "The New Era of AI and its Impact on Data Centres." Technology Magazine. https://technologymagazine.com/articles/the-new-era-of-ai-and-its-impact-on-data-centres

[12] "Power Efficiency and Sustainability." Microsoft Research. https://www.microsoft.com/en-us/research/project/power-capping/?msockid=21545960357564773cd04dc2341a650c

[13] "ChatESG x Lost in Transition(s)." EMEA ESG & Sustainability Research. JP Morgan

[14] Law, Marcus. "The New Era of AI and its Impact on Data Centres." Technology Magazine. https://technologymagazine.com/articles/the-new-era-of-ai-and-its-impact-on-data-centres

[15] Vertiv, Fred R. "What happens when you introduce liquid cooling into an air-cooled data center?" DCD (datacenterdynamics.com). https://www.datacenterdynamics.com/en/opinions/what-happens-when-you-introduce-liquid-cooling-into-an-air-cooled-data-center/

[16] "Team Tackles Thermal Challenge Data Centers Face." NVIDIA Blog. https://blogs.nvidia.com/blog/liquid-cooling-doe-challenge/

[17] "ChatESG x Lost in Transition(s)." EMEA ESG & Sustainability Research. JP Morgan

[18] Berreby, David. "As Use of A.I. Soars, So Does the Energy and Water It Requires." Yale Environment360. Yale School of the Environment. https://e360.yale.edu/features/artificial-intelligence-climate-energy-emissions

Before investing, carefully consider the fund's investment objectives, risks, charges and expenses. Please see the prospectus and summary prospectus containing this and other information which can be obtained by calling 1-866-363-9219. Read it carefully before investing.

Diversification and asset allocation do not guarantee a profit or protect against a loss. Alternative strategies entail added risks and may not be appropriate for all investors. Indexes are unmanaged, not available for direct investment and do not include fees and expenses.

Opinions, estimates, forecasts, and statements of financial market trends that are based on current market conditions constitute our judgment and are subject to change without notice. The views and strategies described may not be appropriate for all investors. References to specific securities, asset classes and financial markets are for illustrative purposes only and are not intended to be, and should not be interpreted as recommendations.

As of 5/31/24, the top 10 holdings of Calamos Antetokounmpo Global Sustainable Equities ETF (% of net assets) include Alphabet, Inc. - Class A 4.6%, Microsoft Corp. 4.5%, Apple, Inc. 3.6%, Taiwan Semiconductor Manufacturing Company, Ltd. (ADR), 3.6%, NVIDIA Corp. 2.3%, Visa, Inc. - Class A 1.7%, SAP, SE 1.6%, Novo Nordisk A/S 1.5%, TJX Companies, Inc. 1.4%, Costco Wholesale Corp. 1.4%.

As of 5/31/24, the top 10 holdings of Calamos Antetokounmpo Sustainable Equities Fund (% of net assets) include Alphabet, Inc. - Class A 7.0%, Microsoft Corp. 6.8%, Apple, Inc. 5.3%, NVIDIA Corp. 3.6%, Taiwan Semiconductor Manufacturing Company, Ltd. (ADR) 3.5%, TJX Companies, Inc. 2.7%, Novo Nordisk A/S (ADR) 2.6%, Merck & Company, Inc. 2.4%, Costco Wholesale Corp. 2.3%, Thermo Fisher Scientific, Inc.

An investment in the Fund(s) is subject to risks, and you could lose money on your investment in the Fund(s). There can be no assurance that the Fund(s)will achieve its investment objective. Your investment in the Fund(s) 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(s) can increase during times of significant market volatility. The Fund(s) also has specific principal risks, which are described below. More detailed information regarding these risks can be found in the Fund's prospectus.

The principal risks of investing in the **Calamos Antetokounmpo Global Sustainable Equities ETF** include: equity securities risk consisting of market prices declining in general, growth stock risk consisting of potential increased volatility due to securities trading at higher multiples, value stock risk, foreign securities risk, forward foreign currency contract risk, emerging markets risk, small and mid-sized company risk and portfolio selection risk. As a result of political or economic instability in foreign countries, there can be special risks associated with investing in foreign securities, including fluctuations in currency exchange rates, increased price volatility and difficulty obtaining information. In addition, emerging markets may present additional risk due to potential for greater economic and political instability in less developed countries.

The principal risks of investing in the **Calamos Antetokounmpo Sustainable Equities Fund** include: equity securities risk consisting of market prices declining in general, growth stock risk consisting of potential increased volatility due to securities trading at higher multiples, large-capitalization stocks as a group could fall out of favor with the market, small and mid-sized company risk, sector risk, portfolio turnover risk, and portfolio selection risk.

The Fund's ESG policy could cause it to perform differently compared to similar funds that do not have such a policy. The application of the social and environmental standards of Calamos Advisors may affect the Fund's exposure to certain issuers, industries, sectors, and factors that may impact the relative financial performance of the Fund-positively or negatively-depending on whether such investments are in or out of favor.

Calamos Antetokounmpo Asset Management LLC ("CGAM"), an investment adviser registered with the SEC under the Investment Advisers Act of 1940, serves as the Fund's adviser ("Adviser"). CGAM is jointly owned by Calamos Advisors LLC and Original C Fund, LLC, an entity whose voting rights are wholly owned by Original PE, LLC which, in turn, is wholly owned by Giannis Sina Ugo Antetokounmpo. Giannis Sina Ugo Antetokounmpo is the majority shareholder of Original C, with a 68% ownership interest.

Mr. Antetokounmpo serves on the Adviser's Board of Directors and has indirect control of half of the Adviser's Board.

Mr. Antetokounmpo is not a portfolio manager of the Fund and will not be involved in the day-to-day management of the Fund's investments, and neither Original C nor Mr. Antetokounmpo shall provide any "investment advice" to the Fund. Mr. Antetokounmpo provided input in selecting the initial strategy for the Fund.

Mr. Antetokounmpo will be involved with marketing efforts on behalf of the Adviser.

If Mr. Antetokounmpo is no longer involved with the Fund or the Adviser then "Antetokounmpo" will be removed from the name of the Fund and the Adviser. Further,shareholders would be notified of any change in the name of the Fund or its strategy.

The Adviser is jointly owned and controlled by Calamos Advisors LLC and, indirectly, by Mr. Antetokounmpo, a well-known professional athlete. Unanticipated events, including,without limitation, death, adverse reputational events or business disputes, could result in Mr. Antetokounmpo no longer being associated or involved with the Adviser. Any such event could adversely impact the Fund and result in shareholders experiencing substantial losses.

NOT FDIC INSURED | MAY LOSE VALUE | NO BANK GUARANTEE

CALAMOS®
INVESTMENTS

Calamos Financial Services LLC, Distributor
2020 Calamos Court | Naperville, IL 60563-2787
866.363.9219 | calamos.com | caminfo@calamos.com

900212 0624