Energy-Efficient AI Networks Lead to Dramatic Reduction in Environmental Impact

Erman Timurdogan
Brian Smith
Peter Roorda

Introduction

Data center hardware infrastructure is undergoing an unprecedented period of expansion, in large part driven by the rapidly growing computational demands of artificial intelligence (AI). Large and complex compute clusters consisting of tens of thousands of processors are today required to meet the needs of generative AI. Due to increasing AI model complexity and the growing use of AI in end-user applications, the demands on data center infrastructure are accelerating. The growth in infrastructure needed for AI is driving a very rapid growth in data center power consumption. Data centers are currently estimated to account for approximately 1% of global energy consumption and CO_2 emissions[1]. As cloud service providers deploy additional and more powerful compute clusters to meet the growing needs of AI applications, emissions and energy consumption are also expected to accelerate, exceeding a 7% CAGR.

Data center infrastructure largely consists of servers and networks. Servers primarily consist of memory, processor, and network interface chips and run software to store, compute, and analyze data. Networks provide high-bandwidth links that distribute data between servers and consist of multiple tiers of electrical switches interconnected by optical transceivers.

In existing cloud data center architectures, networks are responsible for approximately 14% of total power consumption. The emergence of generative AI models, particularly those relying on large language models (LLMs), has led to an exponential rise in both GPU-based server clusters and the switches and transceivers needed for data communications within these clusters given the compute and network intensity of AI compared with other data center applications. The networking needs of AI training has a significant impact on power consumption, and results in AI training networks exceeding 20% of total power consumption. As AI cluster sizes continue to grow, the need for more efficient network solutions becomes even more critical to improve scalability and environmental sustainability.

This paper considers the network power consumption of large-scale AI training implementations and the impact of several emerging approaches to increase power efficiency. This paper highlights the potential of certain key optical innovations to reduce network power consumption by up to 80%, saving more than 17 megawatts of power, and reducing the carbon footprint by the equivalent of 10,000 metric tons of CO_2 per AI training cycle.

AI Training Networks

To tackle the demands of growing AI training workloads, AI training networks increase computational power by connecting numerous compute elements together, including GPUs and TPUs, to enable them to function as a unified supercomputer system. The computational power needed for these large workloads is measured in floating-point operations (FLOPs). The total number of FLOPs is determined by three factors: the total number of parameters (the model's weights), the total number of tokens (training data), and the number of FLOPs required per parameter and per token. The needed FLOPs ultimately dictates the number of processors and accelerators (e.g. GPUs, TPUs, NPUs), and the network capabilities required to support the training cluster.

Our analysis begins by examining the AI data center requirements to train one of the largest AI language models to date, OpenAI's GPT-4[4], and how these requirements change in moving to GPT-5. While numerous AI models exist, these were chosen due to the publicly available data about them. With the network requirements modeled, we compare estimates of network power consumption in moving from GPT-4 to GPT-5, where a significant increase in estimated power consumption and environmental impact is identified.

Our analysis concludes by modeling the use of more energy-efficient optical (EEO) interfaces and utilizing optical circuit switches (OCSs) to replace certain electrical packet switches (EPSs) within the network to significantly reduce the estimated power consumption and environmental impact of GPT-5.

Open AI's GPT-4 was trained over 100 days using 25,000 A100 NVIDIA GPUs. The model contains 1.8 trillion parameters and was trained on a dataset of 13 trillion tokens[5]. The training process for GPT-4 required an estimated 21.5 million exaFLOPs[5-6] or ~2.5 exaFLOPs per second. For comparison, Frontier, the world's fastest supercomputer according to the TOP500 list[7], can only achieve 1.194 exaFLOPs per second.

For GPT-4 training, a non-blocking, multistage network (also known as a folded CLOS or fat tree network) was modeled using EPSs with multiple switch layers[8] to support 25,000 GPUs as depicted in Figure 1. This network utilized 100GPU superpods and 25.6 Tbps electrical packet switches interconnected using 400 Gbps transceivers. To achieve all-to-all connectivity across the network, individual switch radix was insufficient and necessitated the use of multiple switches, as shown in Figure 1.

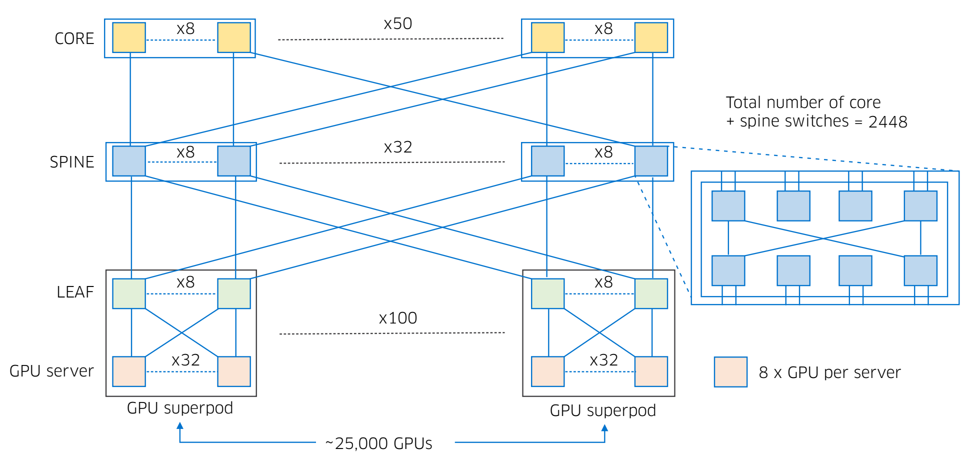

Figure 1. GPT-4 network diagram with 25,000 GPUs and 25.6 Tbps EPS

GPT-5 is expected to consist of 17.5 trillion parameters, a tenfold increase over GPT-4[9]. This mirrors the scaling factor between GPT-3 and GPT-4. For modeling, we assume GPT-5 training will use the same 13 trillion dataset and 100-day training period as GPT-4[10]. Consequently, the compute requirement for GPT-5 is estimated to be an order of magnitude larger than GPT-4, reaching 215 million total exaFLOPs or approximately 25 exaFLOPs per second. This translates to a training system requiring more than 20 times the computing power of the world's current fastest supercomputer.

To estimate the number of GPUs needed for GPT-5 training, we assume next-generation hardware with 2.5 times the performance of the GPUs used in GPT-4 training. Under these assumptions, GPT-5 training would require an estimated 100,000 GPUs. The fat tree network model shown Figure 1 is adapted to support 100,000 GPUs as shown in Figure 2. The adapted model incorporates next-generation electrical packet switches, each with a 51.2 Tbps capacity, interconnected by 800 Gbps transceivers. Where necessary, multiples of these switches are combined to achieve a higher radix. Figures 3a and 3b depict the switch and transceiver counts for GPT-4 and GPT-5 training, respectively.

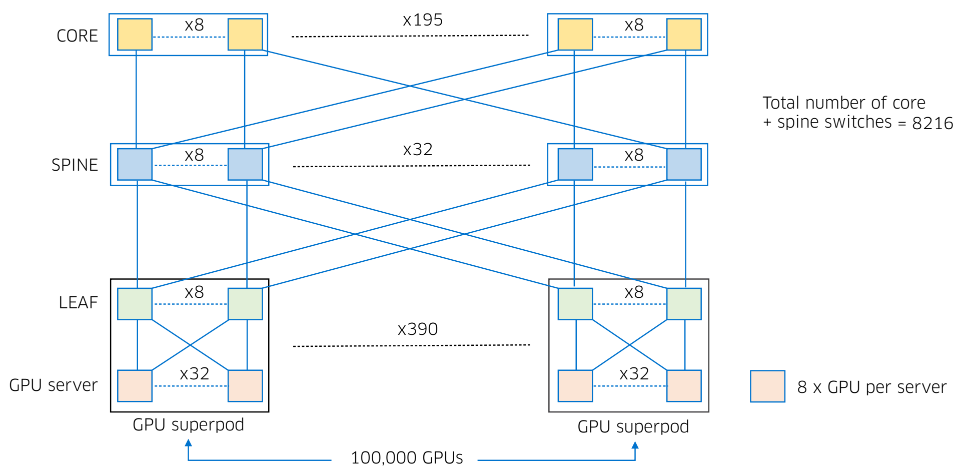

Figure 2. GPT-5 network diagram with 100,000 GPUs and 51.2 Tbps EPS

Figure 3c and 3d show estimates of the total power consumption (including GPUs) and network-only power consumption for GPT-4 and GPT-5 training systems. To put this in perspective, GPT-5 is expected to consume a staggering 122 MW, exceeding 10% of the Hoover Dam's 1,076 MW generation capacity[11]. This stark comparison highlights the critical need for more energy efficient approaches to support the next generation of generative AI models.

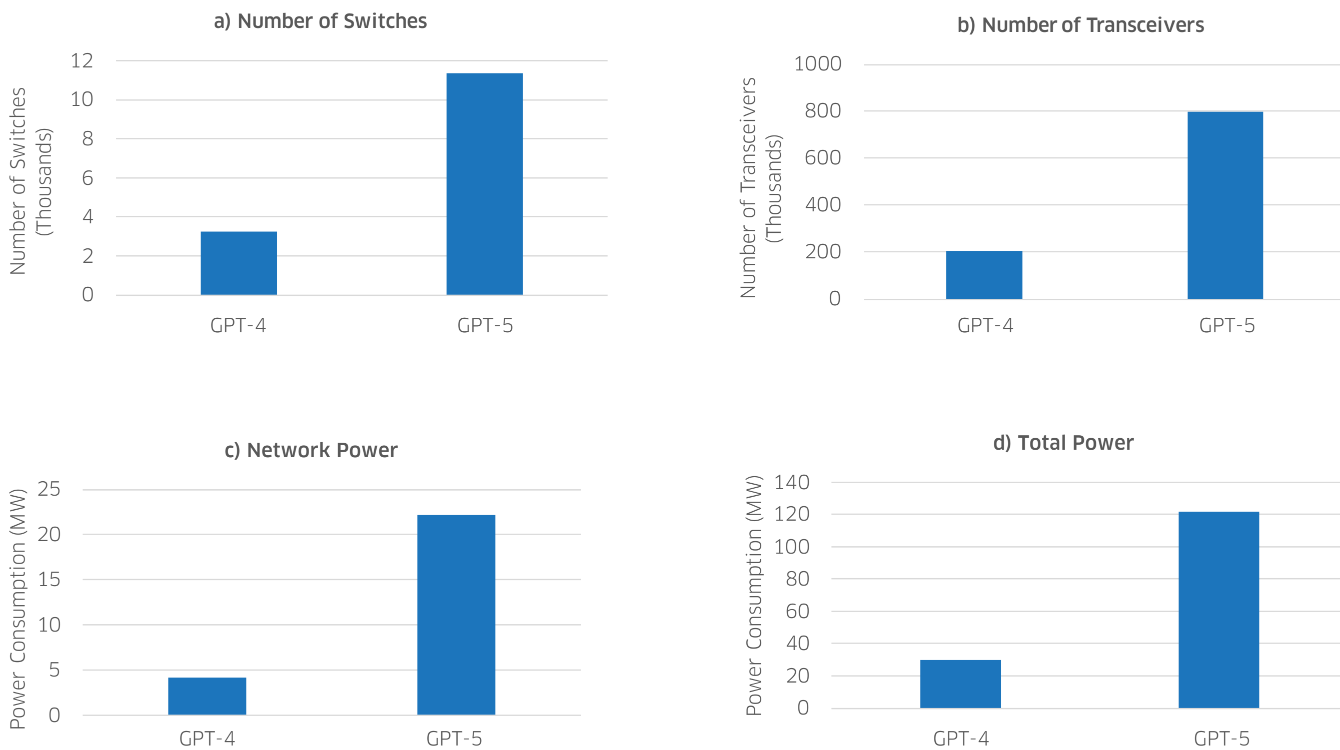

Figure 3. GPT-4 and GPT-5 power and component comparison a) number of EPS needed for each AI model training, b) number of transceivers needed for each AI model training, c) estimated network power comparison between AI models, d) estimated total power comparison between AI models over a single training cycle

Energy-Efficient Networks
In this paper, we focus on two promising areas of innovation that could reduce networking-related power consumption: the utilization of EEO interfaces[12] and OCSs[13] in AI clusters. EEO interfaces can reduce power consumption by eliminating some of the power intensive DSP and/or retiming electronics used in the server-to-switch and switch-to-switch data links. EEO interfaces under development can take various forms, including pluggable transceivers with fewer or no DSPs and/or retimers, or co-packaged optics (CPO) that place optical transceiver functionality in very close proximity to the switch silicon chip. For modeling purposes, we assume EEO interfaces result in an average energy efficiency improvement of approximately 50% for equivalent reach and data rate.

Our modeling shows that incorporating EEO transceivers into the GPT-5 network model can achieve a power savings of 37%, translating to an 8 MW reduction in network power consumption. This saving is significant, nearly double the total network power required for GPT-4 training (Figure 4).

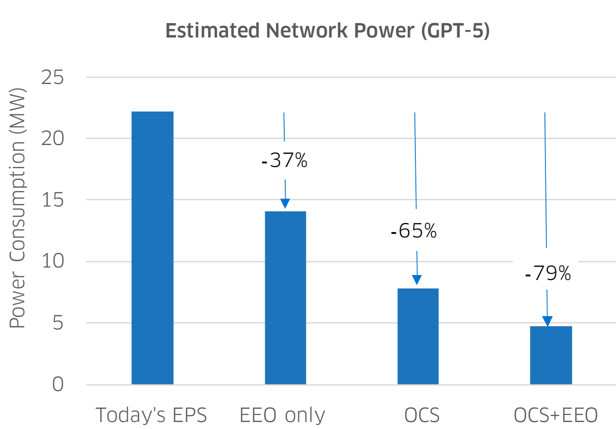

Figure 4. Network power consumption comparison of energy-efficient networks

Utilizing OCS-based high-radix optical switches to replace EPS core network fabric layers allows the flattening of the multitier EPS architecture into a single OCS tier (Figure 5). An OCS can be implemented with a range of underlying technologies, but the most common relies on low-power, electrostatically actuated MEMS mirror technology to direct light from input to output. While this new network architecture requires new software and hardware orchestration, it eliminates the power consumption and latency associated with the large number of electrical switches. EPSs can still be used with short-reach optical interconnects to connect to various network layers. An OCS-based network might require interconnects with longer reach (e.g. FR/LR) optical transceivers to overcome optical switch losses and provide additional reach for connections across the network. This need for longer reach transceiver results in an estimated 20% additional energy consumption per transceiver compared with shorter reach transceivers, but is much more than offset by the power savings of utilizing OCSs. For a network like GPT-5's, a high-radix OCS can connect all GPU superpods within a single switch layer. A unique advantage of utilizing an OCS-based network layer is the ability to expand to larger clusters without altering the network fabric, unlike EPSs which have pre-defined pluggable ports and limited aggregate traffic bandwidth. OCS connections are fully optically transparent and lack the line rate limitations of an EPS-based network.

Moving to an OCS-based network in the GPT-5 example achieves a significant power savings exceeding 14 MW or 65% of total network power, as shown in Figure 4. Additional power savings could be realized by incorporating EEO transceivers. Compared to the baseline network with electrical packet switching, an architecture that combines the benefits of both OCS and EEO technologies could deliver nearly 80% network power savings, translating to a reduction of 17 MW. Utilizing both OCS and EEO, the network power consumption of a GPT-5 system could be comparable to that of GPT-4, while providing an AI training capability an order of magnitude larger.

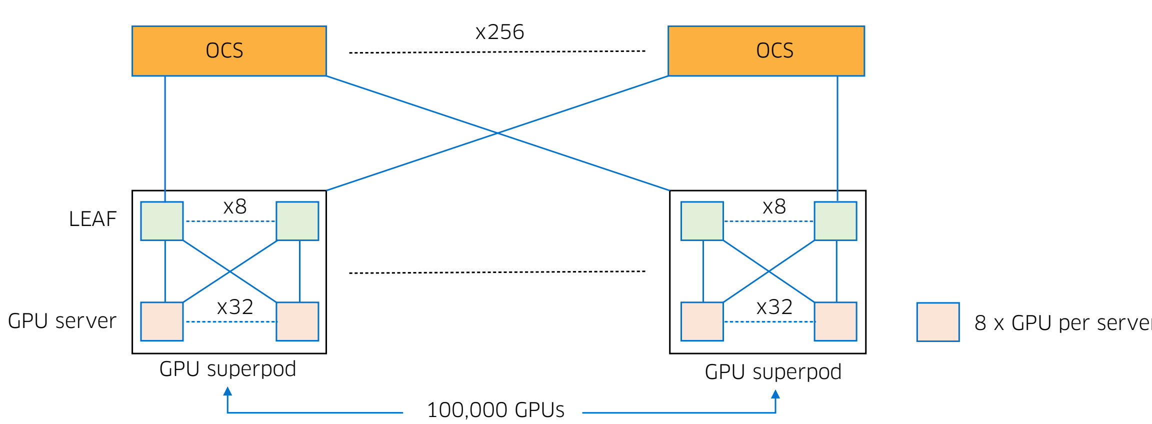

Figure 5. GPT-5 network fabric reconfigured with OCS

The modeled power savings translates to significant reductions in environmental impact. To estimate the CO_2 equivalent (CO_2e) savings from these power reductions, we used a representative carbon intensity metric of 0.24 kg/KWhr, which aligns with that of the Azure West hyperscale data center[14]. Based on this assumption, data center CO_2e savings would exceed 5,000 and 10,000 metric tons for GPT-5 implementations with EEO only and EEO combined with OCS, respectively (Figure 6). To put this into context, the combined CO_2e savings are equivalent to eliminating the emissions from over 2,000 cars in a single year[15].

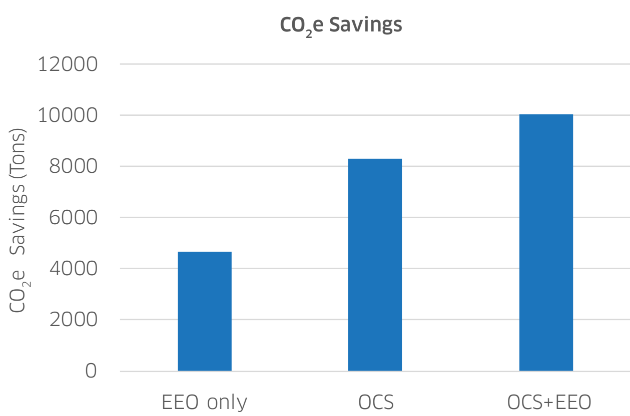

Figure 6. CO_2 emission savings with energy efficient networks

Given the potential of OCS based networks and EEO interfaces to significantly reduce the network power consumption of large AI clusters, we believe the data center industry needs to accelerate efforts into the development and deployment of these technologies to sustainably keep pace with the growing demands of AI workloads. This also requires close collaboration between technology suppliers and hyperscale cloud operators to integrate hardware innovations with software orchestration.

The energy and CO_2 emissions savings that could be achieved with these advancements are substantial. There are also opportunities for even greater efficiency, as the improved latencies offered by combined OCS and EEO networks can lead to better utilization of compute resources and therefore less hardware to achieve a given level of compute performance. Furthermore, deploying the OCS layer within superpods can further reduce power consumption and increase dynamic access to memory. As new AI applications emerge and grow, the benefits of the emerging optical technologies discussed in this paper are even more important to scale data center infrastructure sustainably.

References:
[1] https://www.iea.org/energy-system/buildings/data-centres-and-data-transmission-networks
[2] Cisco Global Cloud Index: Forecast and Methodology, 2011–2016, Cisco White Paper, 2011.
[3] Shi, "Towards Power-Efficient Data-Center Networks", Tu4A.1, OFC 2021
[4] https://openai.com/research/gpt-4
[5] https://klu.ai/blog/gpt-4-llm
[6] https://patmcguinness.substack.com/p/gpt-4-details-revealed
[7] https://www.top500.org/lists/top500/2023/06/
[8] https://docs.nvidia.com/https:/docs.nvidia.com/dgx-superpod-reference-architecture-dgx-h100.pdf
[9] https://dataconomy.com/2023/04/03/chat-gpt5-release-date-agi-meaning-features/
[10] https://cobusgreyling.medium.com/what-are-realistic-gpt-4-size-expectations-73f00c39b832
[11] https://www.cnn.com/2022/08/16/us/hoover-dam-hydropower-drought-climate/index.html
[12] https://www.oiforum.com/oif-launches-forward-looking-project-to-explore-the-next-generation-of-energy-efficient-electrical-and-optical-interfaces/
[13] Poutievski, Leon, et al. "Jupiter evolving: transforming google's datacenter network via optical circuit switches and software-defined networking." Proceedings of the ACM SIGCOMM 2022 Conference. 2022.
[14] https://tinyml.substack.com/p/the-carbon-impact-of-large-language
[15] https://www.epa.gov/greenvehicles/greenhouse-gas-emissions-typical-passenger-vehicle

LUMENTUM

North America
Toll Free: 844 810 LITE (5483)

Outside North America
Toll Free: 800 000 LITE (5483)

China
Toll Free: 400 120 LITE (5483)

eeainetwork-wp-oc-ae 30179715 000 0424