AI DATA CENTERS TIME TO POWER

Ihab Chaaban, P.Eng., MBA
Kassy Hart, M. Sc.
Andrea De Simone, M. Sc.

CONTENTS

INTRODUCTION

Data centers have evolved from relatively predictable, enterprise-focused IT facilities into mission-critical digital infrastructure underpinning cloud computing, hyperscale internet services, cryptocurrency processing, and now generative artificial intelligence (AI) and machine learning (ML). This transformation has changed not only the commercial scale of the sector, but also its electrical and mechanical profile. Facilities that once operated at modest rack densities and relatively stable computational demand now face rapidly shifting loads, significantly higher power densities, and more demanding cooling requirements.

The rise of generative AI has accelerated this shift. AI training clusters and high-performance computing environments can create steep power ramps, pulsed load behavior, high harmonic sensitivity, and stringent uptime requirements. These effects expose weaknesses in both traditional utility-supplied architectures and conventional backup strategies. Data center developers are increasingly constrained by delays in grid interconnection, scarcity of skilled labor, engineering-procurement-construction (EPC) bottlenecks, capital cost escalation, and pressure to deploy capacity faster than utilities can serve it.

In this environment, behind-the-meter (BTM) and island-capable power solutions and time to power (TTP) are becoming strategically important. Among these, aeroderivative gas turbines can offer a compelling option because of their modularity which can help reduce installation time, high power density, fast start capability, operational flexibility, and compatibility with hybrid microgrid architectures. When integrated with battery energy storage systems (BESS), supercapacitors, STATCOMs, advanced controls, and energy management systems (EMS), they can help support stable operation under dynamic AI-related transients while also bridging the gap between demand growth and utility timelines.

This white paper examines the modern data center power challenge from an end-to-end perspective. It begins with the evolution of data centers and the constraints they now face, then explores the role of aeroderivative gas turbines, as an example, in idealized power architectures for islanded and grid-interactive operation. It reviews prior progress in thermal hybrid systems, and torsional vibration risk mitigation. It also discusses the growing importance of inverter-based resources and sustainability pathways. The paper concludes with best practices for closing the power and timeline gap while maintaining reliability, affordability, and system stability.

1. The Evolution of Data Centers and Today's Core Challenges

1.1 From Enterprise Rooms to AI Factories

The data center segment has passed through several distinct phases starting from the Enterprise IT era of the 1990s to mid-2000s where facilities were often built for internal business systems, with moderate power demand, conservative redundancy, and comparatively low (1-3 kW) rack densities.

Fast forward in the timeline, co-location and cloud expansion started to emerge with larger shared facilities, driven by virtualization, internet services, and public cloud providers. Standardization improved, but utility dependence remained dominant.

Then the Hyperscale era with Gigawatt-scale campuses began to appear, often developed by a small number of very large tech companies. These sites emphasized repeatability, modular construction, and energy efficiency.

And then finally, generative AI/ML came along. According to AFCOM's 2026 State of the Data Center Report, the current industry average is 27kW for new rack deployments. Data center campuses demand faster infrastructure turnover, increased liquid cooling adoption and highly dynamic workloads associated with AI training and inference. This final stage is redefining electrical design assumptions. Traditional load profiles for data halls may no longer be sufficient proxies for AI clusters. Operators increasingly must plan for higher instantaneous loads, clustered ramping behavior, and cooling loads that scale alongside compute intensity.

1.2 Growth in Capacity, Rack Density, and Cooling Intensity

Data center loads are rising at three levels simultaneously: 1) Campus scale, from tens of megawatts to hundreds of megawatts; 2) hall scale, with more concentrated power distribution and higher resiliency expectations; and 3) rack scale, where AI racks can far exceed legacy enterprise densities, often requiring direct-to-chip liquid cooling or other advanced thermal management approaches. Hence, higher rack density creates a compounding challenge in terms of more electrical demand per square meter, more heat rejection per rack, greater cooling parasitic load and higher sensitivity to transient events and power quality issues.

Thus, power generation and Power Utilization Effectiveness (PUE) as well as cooling Water Utilization Effectiveness (WUE) are no longer separable design problems. The choice of prime mover affects not only electrical architecture, but also heat integration, black-start capability, resilience, and the economics of rapid deployment.

1.3 Key Challenges Facing Data Centers Today

New, planned data centers face several logistical challenges before they can utilize the compute of their first GPU. The following summarizes key themes seen across industry as more data center projects are announced.

A. Grid power availability

In many regions, the primary constraint on data center development nowadays is the time required to secure power for the facility build-out. Utilities may need several years to deliver upgraded transmission, substations, or firm capacity for such large load (100MW+) consumers. Even where nominal grid capacity may exist, projects can still face energization delays due to transmission bottlenecks, reliability standards, and concerns about overall generation adequacy. As a result, project timelines are increasingly affected by long interconnection queues, network congestion, and the risk of curtailment.

B. Skilled labor shortage

The rapid expansion of data center construction is stressing the skilled labor market across the full project lifecycle, from power development and electrical installation to mechanical systems, controls, commissioning, and facility operations. The demand has accelerated the need for electrical, mechanical, and civil engineers, as well as technicians and tradespeople such as electricians, welders, controls & automation architects, HVAC / liquid cooling experts, and site commissioning specialists. As projects begin to compete for talent, this labor shortage is impacting both schedule and quality.

C. EPC and supply chain constraints

The EPC ecosystem is under significant pressure due to concurrent growth in data centers, grid modernization, industrial electrification and LNG and gas infrastructure. With so much growth in complex projects in parallel, EPCs face longer execution times, additional contingencies to overcome still unknown issues which may arise as data center needs advance.

D. Supply chain and affordability constraints

The surge in data center capacity build out is causing supply chain bottlenecks, leading to extended lead times for critical power assets such as transformers, switchgears, gas turbines, generators, semiconductors and power conversion equipment. Beyond the CAPEX of equipment, which has risen due to demand, total project expenditures are also burdened by the needs for front-of-the-meter (FOTM) interconnection and grid upgrades, as well as fuel, fiber, and water infrastructure. Consequently, selecting technology based solely on upfront costs can be shortsighted; what may appear cheaper at the equipment level may be more expensive at the project level if it delays time-to-revenue.

For many AI projects, TTP is now equal to or more important than levelized cost. If a generation solution allows a facility to begin revenue-generating operation 12 to 36 months earlier than grid service, that schedule advantage can dominate economics. In short, the challenge is not simply how to obtain electricity, but how to obtain firm, scalable, high-quality electricity on a commercially relevant timeline. The chart below shows the intersection between capex and anticipated revenue to illustrate the correlation between these two parameters function of TTP.

Figure 1 examines a 500MW data center operating under a GPU-as-a-Service (GPUaaS) business model, in which tenants are billed based on GPU compute usage per hour. In this model, time to power is a critical economic variable: any delay in energizing the facility postpones operations and increases the opportunity cost of foregone revenue. As grid interconnection timelines lengthen, the revenue lost during the wait can outweigh the upfront capital required to deploy on-site or dedicated power generation. Beyond a certain threshold, it may be more economical to invest early in alternative power infrastructure and begin monetizing capacity sooner, rather than wait several years for grid access.

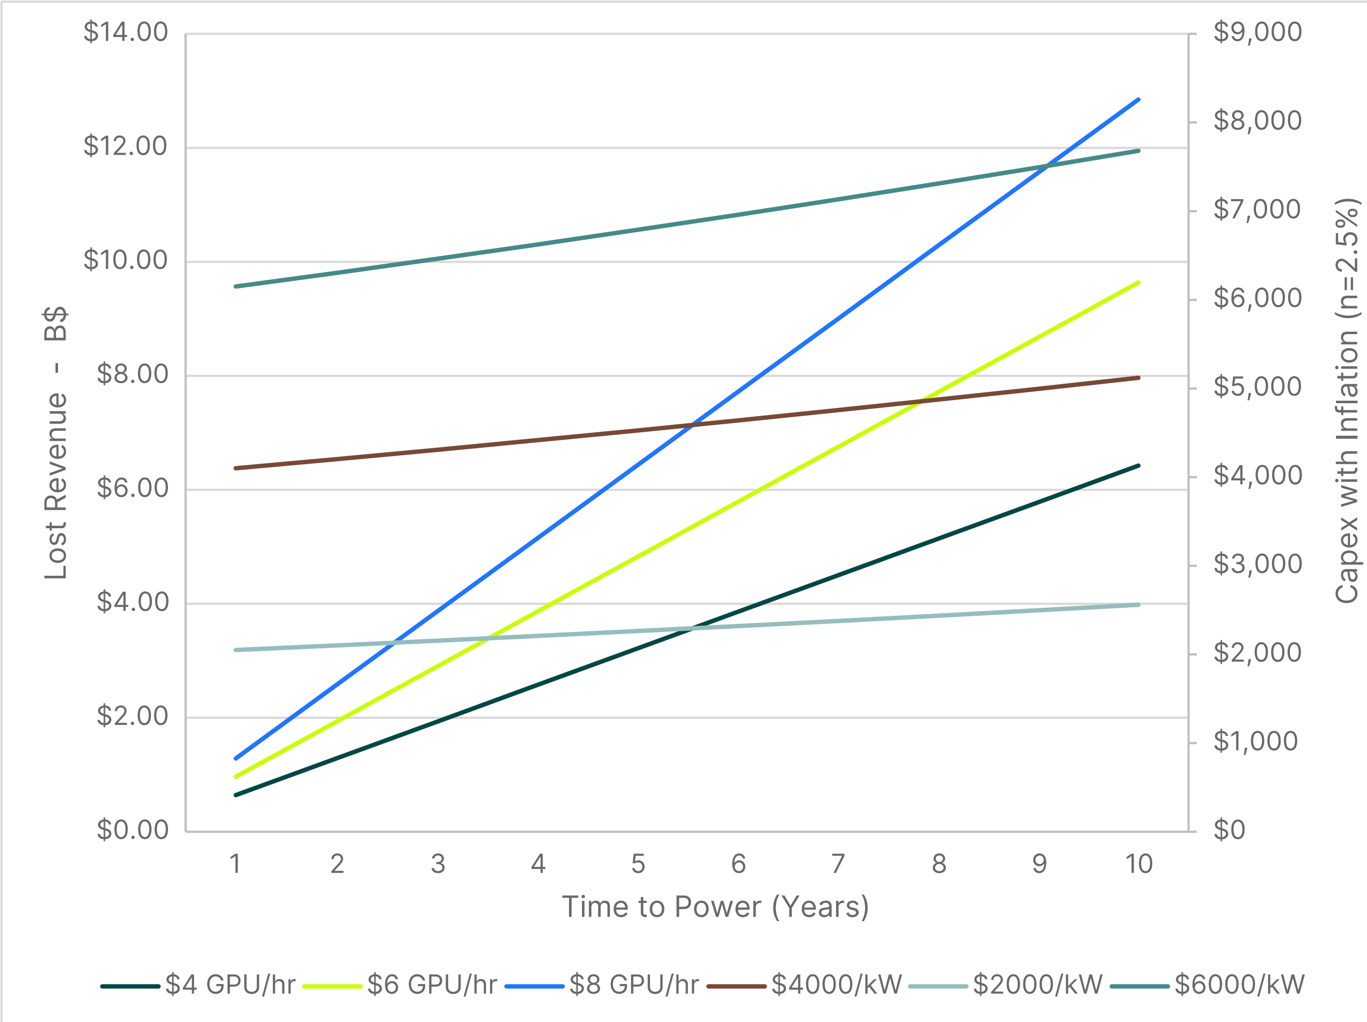

*Figure 1: Project Viability

*The performance results, economic analysis, and hybrid system examples presented in this paper are illustrative and based on representative assumptions. Actual project performance, economics, timelines, emissions, and system configurations will vary depending on site conditions, operating profiles, equipment selection, market conditions, and applicable regulatory requirements.

2. The Data Center Power Model and the Role of Aeroderivative Gas Turbines

2.1 What an effective power model for a data center looks like

A power architecture well-suited for next-generation data centers is a highly integrated, grid-interactive system designed to enable rapid deployment while delivering high availability, scalability, and N+2 or greater (see Table 1 for example) resiliency consistent with Tier III and higher facility requirements. The model can provide black-start capability, stable transient performance under dynamic operating conditions, and islanding functionality to maintain continuity during grid disturbances or outages. It must also be engineered to support the fast-ramping, highly variable load profiles associated with AI/ML workloads, while preserving power quality and system stability. System architecture needs to maintain flexibility for future decarbonization through optional integration of low-carbon generation, storage, and other emerging energy resources.

From an infrastructure standpoint, one practical approach is typically not based on a single technology but on a hybrid, layered architecture that combines firm dispatchable generation, fast-acting power electronics, and energy storage for energy shifting, backup power, and transient damping. This architecture should be coordinated through an advanced energy management system (EMS) and protection scheme, with integrated synchronization and islanding controls to maintain system stability and operational continuity. It may also incorporate cooling-load coordination to optimize overall facility performance and, where feasible, preserve future fuel flexibility to support evolving operational and decarbonization requirements.

2.2 Why Aeroderivative Gas Turbines fit this Model

Aeroderivative gas turbines are derived from aviation engine technology and are well suited for applications where speed, modularity, and operational flexibility matter. Compared with many other thermal assets, they offer several product attributes attractive to large data centers, as described below.

A. High power density

Aeroderivative gas turbines offer high power density compared with other gas-fired generation technologies. Even before accounting for Balance of Plant equipment – which is required across all power generation solutions - Table 2 below highlights the advantage of Aeroderivative gas turbines.

Technology	Space Requirement (sq-ft/MW)
Heavy Duty CCGT	1000 – 4000
Heavy Duty SCGT	500 – 2000
Aeroderivative GT	340 – 520
Recips – Medium Speed	1150 – 1800
Recips – High Speed	550 – 1800

Table 2: Comparison of different gas power generation technology by footprint requirements

In addition, because of their sizes, the Aeroderivative gas turbines offer a prime solution to avoid the overbuilt capacity, especially when the need for redundancy is high. Figure 2 shows the response to certain power needs of different block sized solutions, illustrating how much overcapacity must be considered when a certain reliability or availability requirement must be met (i.e. 99.999%, also called “5 × 9s”) to meet industry standards such as Uptime Institute Tier III/IV or TIA-942. When a big 500 MW block solution is considered, depending on the reliability of the equipment, smaller engines can require double the installed base needed, whereas Aeroderivative gas turbines might require as little as just a modest 5 to 10% increase in MW ordered.

Availability Target	Product Availability																							
	87.50%	88.00%	88.50%	89.00%	89.50%	90.00%	90.50%	91.00%	91.50%	92.00%	92.50%	93.00%	93.50%	94.00%	94.50%	95.00%	95.50%	96.00%	96.50%	97.00%	97.50%	98.00%	98.50%	99.00%
99.0000%	3	3	3	3	3	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	1
99.9000%	4	4	4	4	4	3	3	3	3	3	3	3	3	3	3	3	3	3	3	2	2	2	2	2
99.9900%	5	5	5	5	5	5	4	4	4	4	4	4	4	4	4	4	3	3	3	3	3	3	3	2
99.9990%	6	6	6	6	6	6	5	5	5	5	5	5	5	5	4	4	4	4	4	4	4	3	3	3

Table 1: Level of redundancy needed as function of product availability and required overall power generation availability

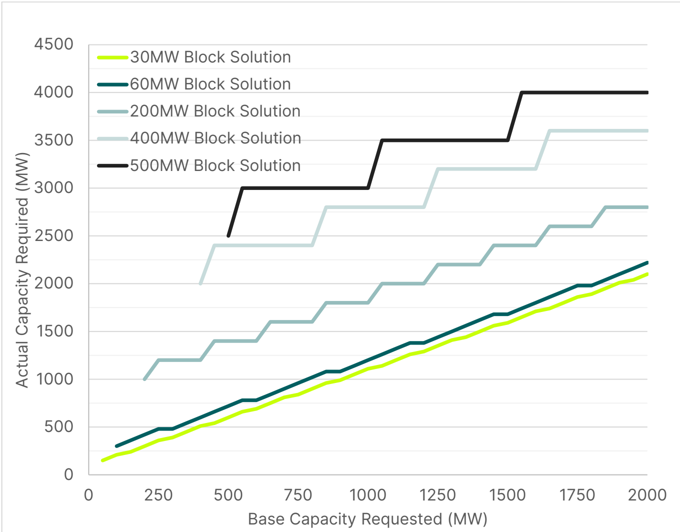

*Figure 2: Power requested versus Power needed for a 99.999% (5 × 9's) availability for different block sizes

B. Rapid-start capability

BESS are excellent for instantaneous response, but their duration is limited. Special provisions have been engineered for the Aeroderivative gas turbine engine and its package auxiliaries to accommodate fast start capability as low as 5 minutes from cold conditions for certain models. This rapid-start capability can bridge from battery support to sustained generation more effectively, reducing the amount of storage required for contingency coverage.

C. Modular deployment

Data center buildouts are increasingly executed in phases. The modular nature of the Aeroderivatives, specifically the GE Vernova Gas Power LM2500XPRESS and TM2500 packages, allow for scalable power density. An operator can deploy a single unit to power the first phase of a facility or its construction and add additional modular units in manageable blocks as the IT load grows.

The LM2500XPRESS package has been engineered to move the assembly process from the field to a controlled factory environment. Where traditional power plant construction is built at site, requiring multitudes of parts to be shipped to site for assembly over many months, the LM2500XPRESS is modularized into pre-tested, transportable skids, which can be shipped ahead of the gas turbine engine. By the time the turbine arrives at the data center site, a vast majority of the piping, wiring, and instrumentation is already connected and commissioned at the factory. This can significantly reduce on-site construction time, labor requirements, and the risk of integration errors.

Figure 3: GE Vernova Gas Power's TM2500 - modular packaging with mobility enabled.

The TM2500 (see Figure 3) holds a unique competitive advantage for data center developers due to its trailer-mounted, fully mobile nature. If a data center's power needs are temporary (e.g., waiting for a local utility substation to be upgraded) or if a facility needs to shift capacity to a different site location, the TM2500 units can be disconnected, hauled away, and redeployed elsewhere in weeks. As the TM2500 is designed to operate on a simple concrete pad rather than requiring the intensive foundation work of a permanent plant, it can also reduce the civil engineering burden of the project.

GE Vernova's LM2500 portfolio shifts the paradigm of power generation from a civil construction project to a standardized equipment deployment.

D. High simple-cycle efficiency

With their relatively high efficiency in their class, aeroderivative gas turbines can achieve simple-cycle efficiencies of 38 to 42% (lower heating value) subject to the configuration and site conditions.

E. Operational flexibility

Aeroderivative gas turbines are derived from jet engine technology, making them well-suited for operational peaking, cycling, and island service. For data centers, this makes Aeroderivatives a compelling backbone in a behind-the-meter or hybrid microgrid context. Additionally, depending on design and fuel strategy, Aeroderivative gas turbines have potential for compatibility with lower-carbon fuels, such as e-fuels, bio-mass and hydrogen, providing not only flexibility in conventional gas and liquid fuels but also a pathway for decarbonization.

2.3 Other technologies compared to Aeroderivatives

It should be emphasized that Aeroderivative technology is not subject to the ISO 8528 step-load performance standard typically applicable to reciprocating engines. The mentioned standard (whether it is class G1, G2 or G3) describes how a reciprocating engine-generator set can accept a sudden increase in electrical load while maintaining stable frequency

*The performance results, economic analysis, and hybrid system examples presented in this paper are illustrative and based on representative assumptions. Actual project performance, economics, timelines, emissions, and system configurations will vary depending on site conditions, operating profiles, equipment selection, market conditions, and applicable regulatory requirements.

and voltage within preset ranges, while brake mean effective pressure (BMEP) helps explain the engine capability behind that response. In conventional applications, when a large load is applied in steps, as guided by the code, the reciprocating engine must quickly produce more torque to prevent excessive speed drop and recover within the limits defined by ISO 8528. That torque demand is reflected in BMEP, which is a standard way of expressing how hard the engine is working relative to its size.

On the other hand, the Aeroderivatives are two- and three-shaft type designs with control systems that have very robust fuel control acceleration and deceleration schedules that can maintain stall margins suitable for the applications. The Aero high-pressure shaft is aerodynamically coupled to the low-pressure shaft; therefore, the AC generator frequency output is controlled via the speed of the low-pressure turbine.

2.4 Functional Role of Aeroderivatives in a Hybrid Data Center Power System

In hybrid data center power systems serving AI/ML-driven dynamic load environments, Aeroderivative gas turbines can provide significant value as part of an integrated architecture rather than as standalone assets. Within this framework, the Aeroderivative provides bulk firm power and synchronous inertia, while complementary power-conversion resources—including battery energy storage systems (BESS), supercapacitors, STATCOMs or other forms of dynamic reactive power support and voltage regulation, fast load-shedding, EMS platforms, and informed protection and damping strategies—manage rapid load ramps, sub-second transients, voltage stabilization, and ride-through performance. The EMS or microgrid controller coordinates dispatch, state of charge, spinning reserve, and mode transitions to ensure stable, resilient, and responsive system operation.

2.5 Illustrative Islanded operation challenges

When time to power (TTP) is a key project requirement, generation assets may need to operate in islanded mode for prolonged durations — commonly referred to as behind-the-meter (BTM) operation—until utility grid interconnection is established. This interim operating configuration introduces distinct technical and operational challenges, including load-generation balancing, voltage and frequency regulation, protection coordination, black-start capability, and overall system stability. Accordingly, islanded operation must be carefully engineered, commissioned, and managed to ensure reliable and secure performance. Below are conceptual charts showing the role of a hybrid system (Aeroderivative + BESS) during islanded operation under a typical AI/ML dynamic load. These are illustrative, not site-specific design results.

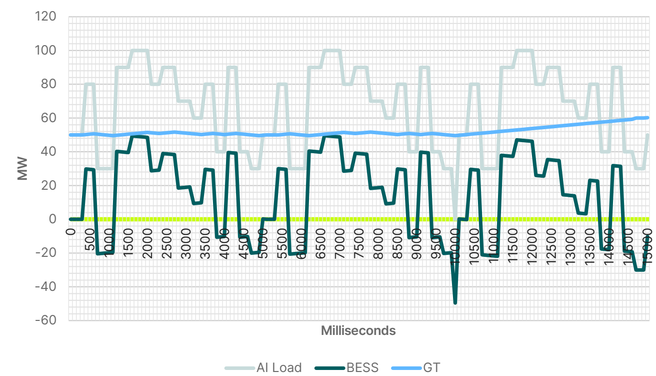

*Figure 4: Load-BESS-Gas Turbine Operation - typical AI load curve and associated hybrid operation between an Aero GT and BESS

The highly dynamic behavior of AI-related computational demand, as shown in Figure 4, introduces significant challenges due to rapid temporal variations in data center power consumption. Such load behavior may be represented through a user-defined, time-varying current or power waveform constructed from a measured or synthesized time-series load profile. From a signal-processing perspective, the load profile can be decomposed into a set of constituent components that collectively describe its temporal characteristics and corresponding spectral content. By applying Fourier transform methods to the load profile, it is possible to quantify the frequency-domain representation of the load and identify the frequency bands with the highest energy concentration within the torsional frequency range (see Figure 4 and Figure 5). This analysis enables characterization of the dominant frequency components that may influence system dynamics, power quality, and network response.

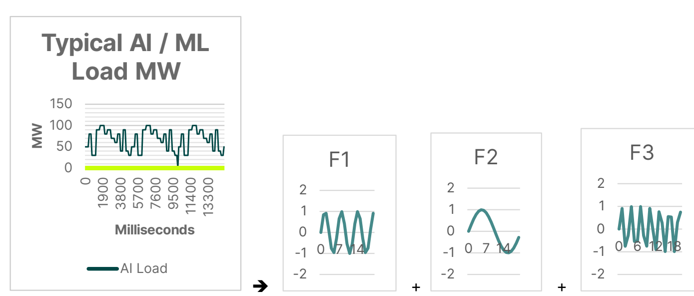

Figure 5: *Deconstructing the load curve

*The performance results, economic analysis, and hybrid system examples presented in this paper are illustrative and based on representative assumptions. Actual project performance, economics, timelines, emissions, and system configurations will vary depending on site conditions, operating profiles, equipment selection, market conditions, and applicable regulatory requirements.

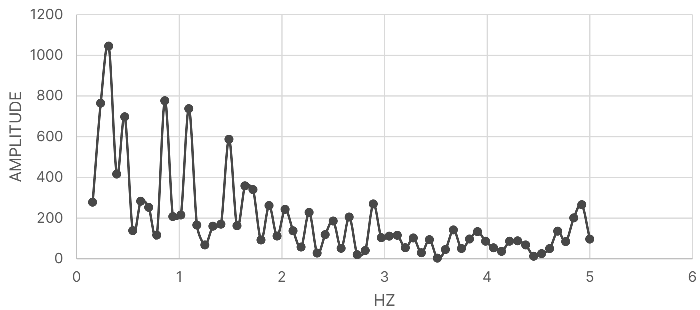

*Figure 6: Example of an FFT to analyze the frequency spectrum in the load curve – for illustration purposes only

2.6 Torsional vibrations

Torsional vibration is the oscillatory angular deformation of rotating shaft systems. Although it is often less apparent than linear vibration, torsional vibration can impose severe mechanical stress and lead to significant equipment damage. In power generation and conversion systems, torsional excitation may be driven by electrical torque pulsations, forced power oscillations, generator–network interactions, converter control interactions, sub synchronous torsional interaction, and resonance within the shaft train. These mechanisms can excite one or more torsional modes of the rotor shaft system, resulting in cyclic torque amplification and the accumulation of fatigue damage. If not properly identified and mitigated, such conditions may cause shaft fatigue, coupling degradation, gear tooth damage, and, in severe cases, failure under repeated stress cycles.

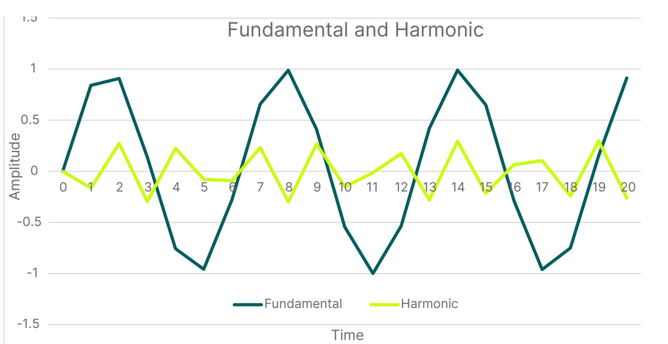

Figure 7: Example of a fundamental and harmonic wave

A white paper co-authored by Microsoft, Open AI, and NVIDIA in 2025 outlines the challenges of power stabilization for AI training data centers. Evaluation of the impacts associated with data center operations requires a broad set of technical studies, each of which depends on models tailored to the specific phenomenon under investigation. Among the various model types required is electromagnetic transient (EMT) modeling. EMT modeling represents one of the highest fidelity approaches available in power system analysis, as it captures fast transients and detailed dynamic interactions with a high degree of temporal resolution. Developing a credible EMT model, however, requires substantial technical expertise as well as access to detailed equipment and system data. The outcome would provide recommendations on the overall architecture of an integrated system that provides necessary damping as well as torsional stress protection and monitoring to ensure stable and reliable operation.

3.Thermal hybrids for speed to power

3.1 Earlier Efforts That Enabled Today's Solutions

The hybrid power asset solution space did not emerge abruptly; rather, it has evolved through decades of experience across multiple power system applications. Its development has been informed by lessons learned from industrial microgrids, where diverse generation and load resources have long been coordinated to support reliability, resilience, and operational flexibility under both grid-connected and islanded conditions. In parallel, thermal hybrid configurations—including GEV's Enhanced Gas Turbine (EGT)–based systems—have advanced as a means of improving dispatchability, dynamic performance, and overall plant efficiency by integrating conventional thermal generation with complementary control, storage, or balance-of-plant technologies. More recently, the rapid deployment of renewable hybrid systems has further expanded the design and control framework for hybrid power plants, particularly through the integration of wind, solar, and battery energy storage in both utility-scale and behind-the-meter applications. Collectively, these historical developments have established the technical foundation for today's hybrid solution space, which leverages proven concepts from thermal generation, industrial power systems, and renewable integration to address increasingly complex performance, resilience, and interconnection requirements.

3.2 Leveraging thermal hybrids for data centers

As data centers face supply chain constraints and increasing pressure to meet decarbonization targets, thermal hybrid architectures represent a viable alternative for onsite power provision. A hybrid configuration combining wind generation, solar photovoltaic (PV) resources, Aeroderivative gas turbines, and BESS can reduce the operating duty and cumulative run hours of the gas turbines by dispatching them primarily to firm variable renewable generation and maintain reliability during renewable shortfalls. Such an approach can help improve the project TTP via the diversity of assets, lower overall carbon emissions, improve fuel utilization, and

*The performance results, economic analysis, and hybrid system examples presented in this paper are illustrative and based on representative assumptions. Actual project performance, economics, timelines, emissions, and system configurations will vary depending on site conditions, operating profiles, equipment selection, market conditions, and applicable regulatory requirements.

potentially support reducing the total installed gas turbine capacity required to meet data center demand. In addition, the inclusion of BESS can enhance fast frequency response, ramp-rate control, and transient performance, thereby improving the operational flexibility and dynamic stability of the hybrid power system.

3.3 A typical thermal hybrid use case

While renewable energy assets offer a pathway to carbon-free electricity generation, their intrinsic intermittency and variability may limit their ability to independently satisfy the continuous and highly reliable power requirements of data center loads. To address this challenge, a hybrid generation architecture was evaluated for a representative site consisting of a 300 MW wind facility, a 300 MW single-axis monocrystalline solar PV installation, and 150 MW of dispatchable firming capacity provided by five LM2500+G4 XPRESS Aeroderivative gas turbine units. This integrated hybrid system was configured to serve a 150 MW behind-the-meter (BTM) data center load with possible consideration to export excess renewable energy to the grid to reduce curtailment.

In this modeled example, analysis indicates that this hybrid configuration can achieve a capacity factor in excess of 95% while delivering a substantial reduction in CO2 emissions relative to a fully gas-fired generation scenario. The inclusion of dispatchable thermal capacity mitigates the intermittency associated with variable renewable resources, thereby enabling a level of supply reliability more closely aligned with data center operational requirements.

For comparison, a conventional co-located wind and solar configuration at a site characterized by strong wind resource and favorable solar global horizontal irradiance (GHI), when sized solely to the point of interconnection (POI) rated capacity, would typically be limited to a maximum capacity factor of approximately 65%. In this context, the proposed hybrid architecture represents a diversified and operationally robust generation model that materially increases annual renewable energy production while offering developers an alternative pathway to achieving both reliability and emissions reduction objectives in data center power applications that is challenged by supply chain constraints.

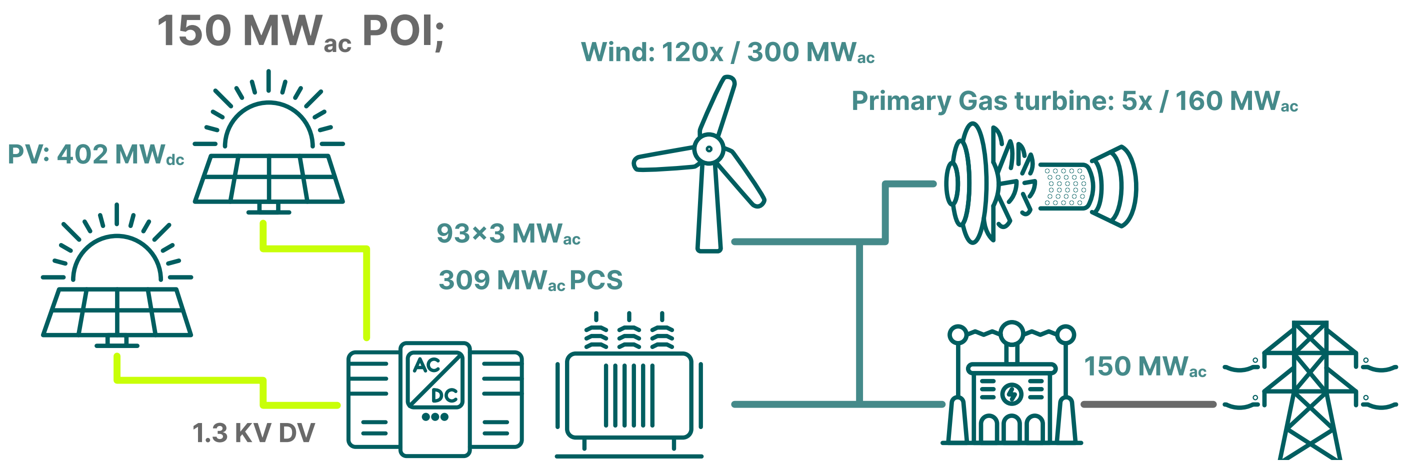

Figure 8: A typical thermal hybrid architecture consisting of a 300 MW Wind farm, 300 MW Solar PV and 160 MW Aeroderivative thermal asset (5 x LM2500XPRESS)

Hours ON	1340.4 hrs.	# Starts	379
% ON	15.30 %	# Cold Starts	28.8
Hours OFF	7419.6 hrs.	# Warm Starts	181
% OFF	84.70 %	# Hot Starts	169.2
Plant Load Factor	11.71 %	Avg. Shutdown duration before Cold Start	121.69 hrs.
Load ON	75.64%	Avg. Shutdown duration before Warm Start	18.56 hrs.
Efficiency ON	33.32 %	Avg. Shutdown duration before Hot Start	3.18 hrs.
Annual CO2 Emission (mTon)	97,021		

Table 3: Year 1 key performance indicators for the thermal assets in the arbitrary chosen use case.

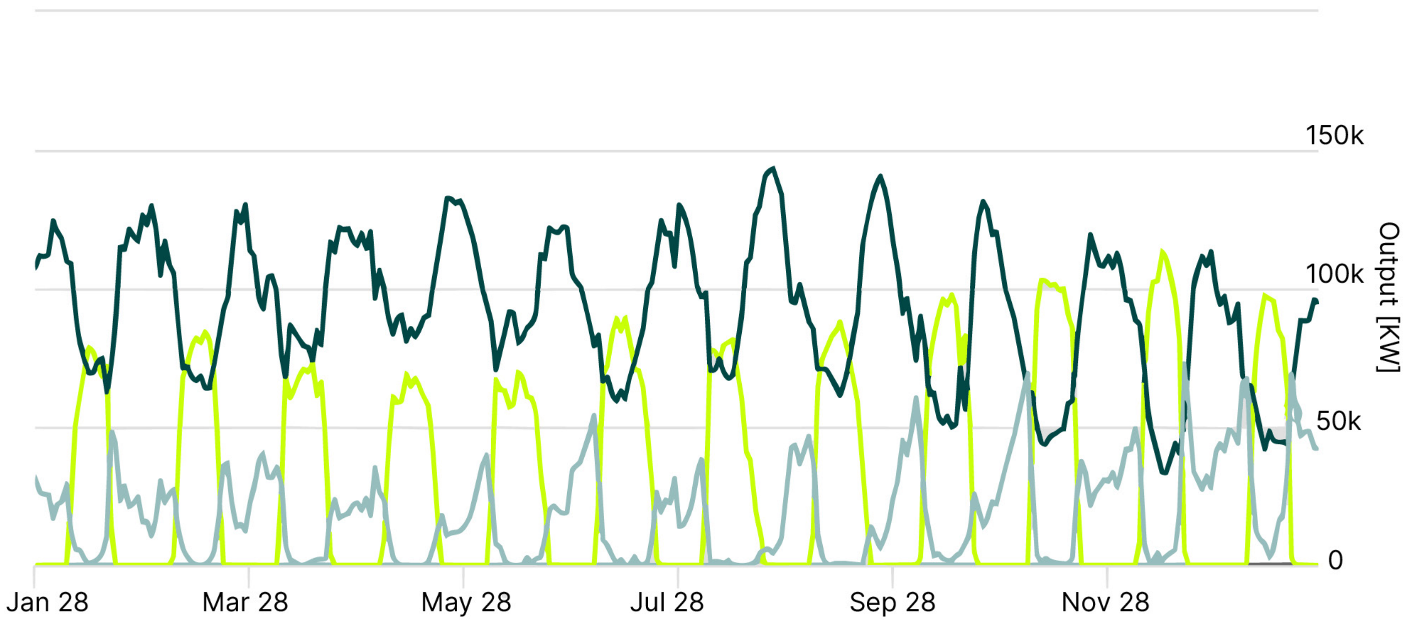

Figure 9: Year 1 hourly output – Blue: Dark Green, Urgency Green: Solar PV, light Green: Aero Gas Turbines

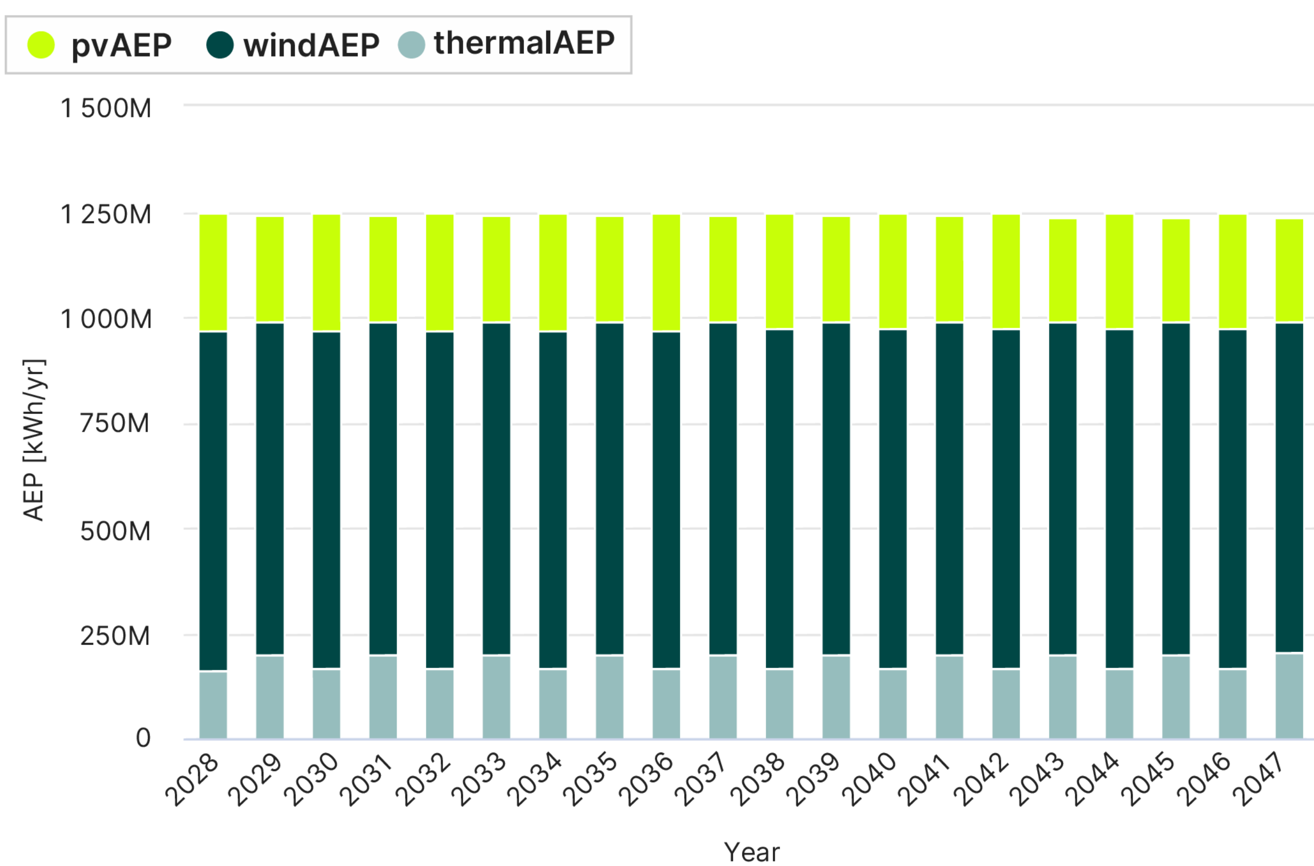

Figure 10: Annual Energy Production per asset

3.4 Role of the Energy Management System

The Energy Management System (EMS) constitutes supervisory control architecture that integrates generation, storage, substation, and load-side assets into a unified and coordinated operating platform. In this context, the EMS framework includes the Master EMS/PMCS SCADA system, real-time power management controllers, the Battery Management System (BMS), protection relays, programmable logic controllers (PLCs), remote terminal unit inputs/outputs (RTUs), network communications, and cybersecurity systems. Collectively, these elements enable real-time monitoring, command and control, protection coordination, improved coordination of assets, and secure system-wide communications across the power ecosystem. For applications supporting AI-driven data center loads, the quality and integration of the EMS are of increasing strategic importance, as the supervisory and control architecture must accommodate highly dynamic demand behavior while maintaining electrical stability, operational resilience, and overall system performance.

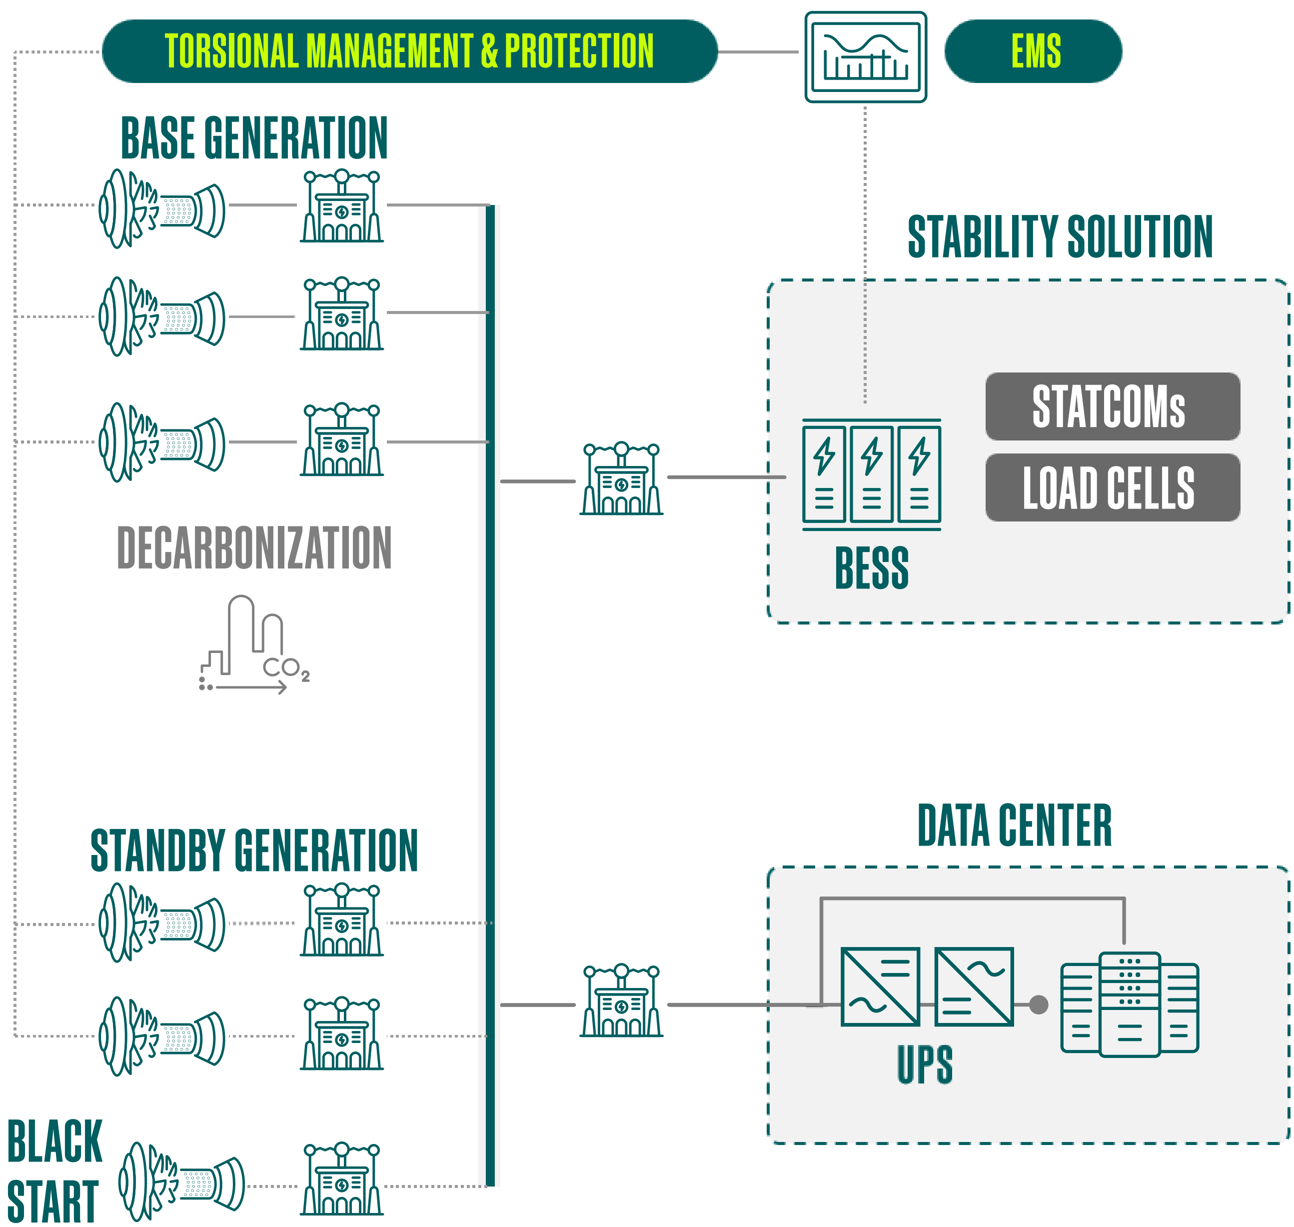

Figure 11: High-level system diagram of Energy Management System for integrated solution

4. Inverter-Based Resources, Short-Circuit Characteristics, and Power Electronics Implications for AI Data Centers

Inverter-based resources (IBRs), including solar PV and battery energy storage systems (BESS), differ fundamentally from conventional synchronous generation in both fault behavior and dynamic system support. Unlike synchronous machines, which inherently contribute high short-circuit current and rotational inertia, IBRs are interfaced through power electronics and are therefore typically current-limited by design. As a result, their fault current contribution is substantially lower in magnitude and shorter in duration than that of synchronous generators. This characteristic has important implications for protection system design, voltage recovery, transient stability, and the overall robustness of power systems supplying AI-oriented data centers.

Synchronous generators provide electromechanical inertia through their rotating mass, which naturally resists rapid changes in system frequency and helps stabilize the power system during disturbances. IBRs, by contrast, do not inherently contribute physical inertia unless specifically programmed to emulate such behavior through advanced control schemes, often referred to as synthetic or virtual inertia. Although these controls can improve system response, they are not a direct substitute for the instantaneous and naturally available inertial response of synchronous machines. Consequently, systems with high IBR concentration may exhibit faster frequency excursions, lower fault current availability, and greater sensitivity to disturbances unless these characteristics are explicitly addressed in the system architecture and proper modeling.

These issues are particularly relevant for AI data centers, where load behavior can be highly dynamic due to rapid compute ramping, accelerator clustering, and power fluctuations associated with advanced processing hardware. In certain architectures, power electronics are introduced between the upstream power system and the IT load to electrically isolate the server racks from fast load disturbances, either through LV or MV systems. This approach can help decouple highly variable rack-level demand from the broader generation and distribution system, thereby reducing the propagation of transient disturbances, improving local power quality, stability, and limiting the impact of sudden load step changes on upstream assets. Such architectures may also enable tighter voltage regulation and improved coordination with energy storage and fast-acting controls.

However, this increased reliance on power electronics introduces additional system considerations. Power electronic interfaces can alter fault signatures, reduce system strength, and create complex dynamic interactions among inverters, converters, protection devices, and control systems. They may also increase sensitivity to control instability, harmonics, and electromagnetic interactions if not properly modeled and coordinated. As a result, architectures that rely heavily on IBRs and power-electronic isolation require more advanced dynamic studies than conventional data center power systems, including electromagnetic transient analysis and detailed protection coordination reviews.

Accordingly, while IBR-based and power-electronics-dominant architectures offer meaningful advantages in flexibility, decarbonization, and load conditioning for AI data centers, they also require a fundamentally different approach to system architecture. Short-circuit performance, inertia deficiency, control coordination, and dynamic load isolation must be addressed as primary engineering considerations to ensure that reliability, resilience, and power quality objectives are achieved under increasingly demanding AI load profiles.

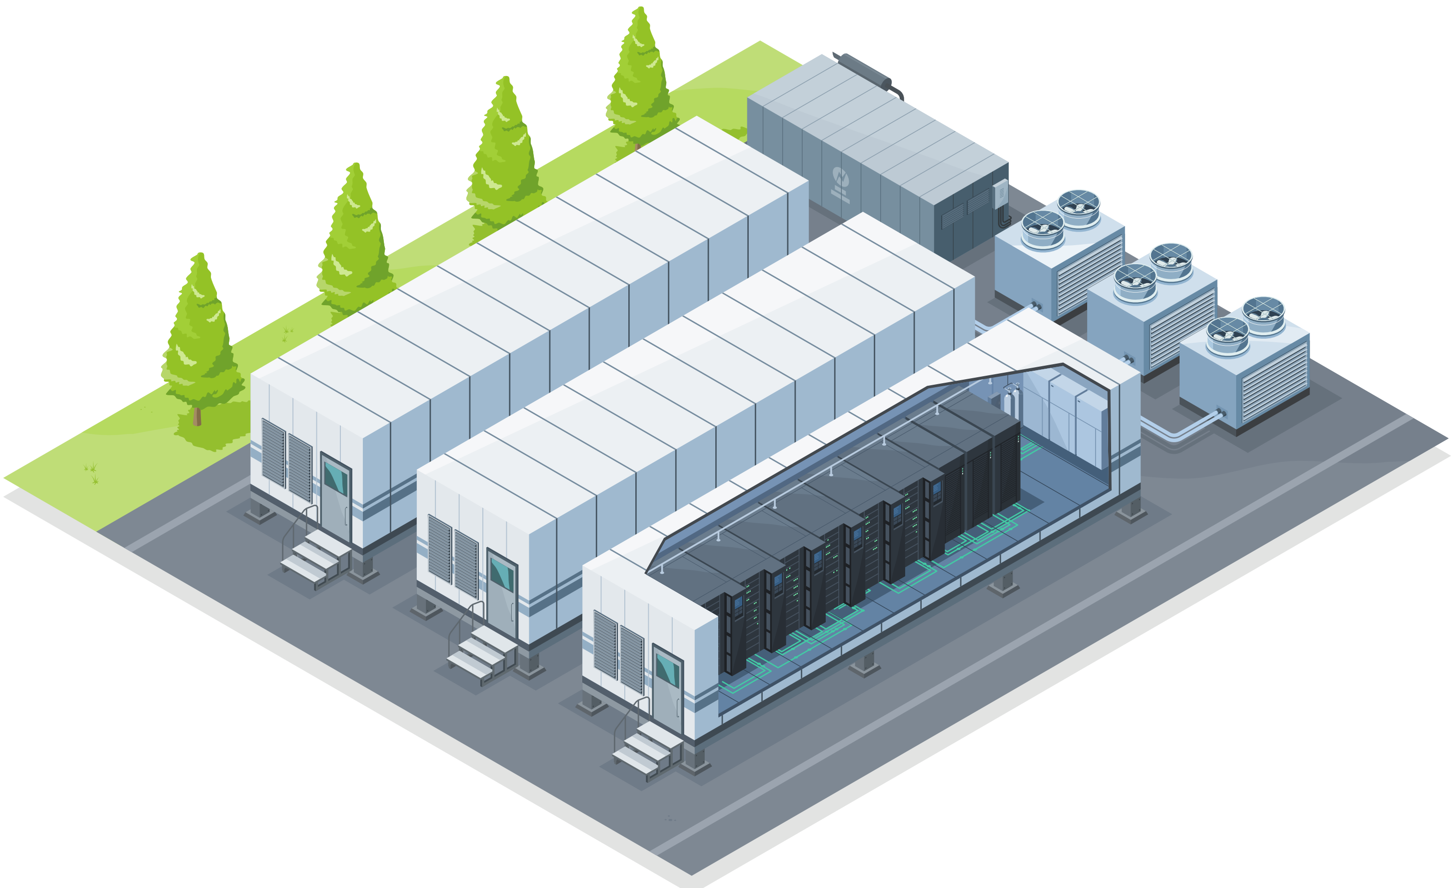

CONCLUSION

Artificial intelligence is driving a step change in digital infrastructure requirements; however, the advancement of AI is increasingly outstripping the readiness of the physical systems needed to support it at scale. While progress in models, compute hardware, and software frameworks remains rapid, the enabling infrastructure stack—including power supply, thermal management, supply chain capacity, labor availability, and EPC execution—has emerged as a primary constraint on deployment. Consequently, the central challenge associated with AI growth is no longer exclusively just a computational hurdle, but an infrastructural, operational, and systemic one.

This shift is particularly evident in data center development, where AI-driven loads are exposing limitations in traditional design approaches. Long grid interconnection timelines, conventional diesel-only backup strategies, low-density electrical configurations, cooling systems optimized for lower heat flux, are often misaligned with the requirements of high-density, highly dynamic AI workloads. These conditions place new emphasis on reliable power delivery, advanced cooling architectures, accelerated deployment models, and system designs capable of maintaining stability and performance under rapid load variation, particularly islanded architectures.

Consequently, infrastructure strategies that were effective for previous generations of data centers may not be sufficient for the AI era. The next phase of development is likely to require an integrated approach in which power, cooling, resilience, power quality, and project execution are designed as interdependent elements of a single operating architecture. In this environment, competitive advantage will increasingly depend not only on access to compute technology, but also on the ability to deliver robust, scalable, and deployment-ready infrastructure capable of supporting AI at industrial scale.

One practical path forward is not ideological attachment to a single technology. It is a practical, technically rigorous hybrid approach that combines the durable benefits of synchronous generation with the speed and flexibility of inverter-based resources. That architecture can bridge the gap between demand growth and grid readiness, while supporting the reliability, affordability, and stability required by the next era of digital infrastructure.

6. APPENDIX

6.1 References

1. Modelling and Simulation of Gas Turbine and Controlling the Sudden Change in Load
   Abdulgani Albagul1, Montaser .A. Qasem2
2. Pacific Northwest National Laboratory – Electromagnetic Transient Modeling of Large Data Centers for Grid-Level Studies Alpha Release - December 2025 - Brett A. Ross , Jim Follum
3. NERC (North American Electric Reliability Corporation) – Inverter-Based Resource Performance and Reliability Guidance **https://www.nerc.com**
4. GE Vernova – Aeroderivative Gas Turbine Product and Technology Information
   https://www.gevernova.com/gas-power/products/gas-turbines/Aeroderivative-gas-turbines
5. GE Vernova – Hybrid Power, Grid Stability, and Power Conversion Solutions
   https://www.gevernova.com
6. National Renewable Energy Laboratory (NREL) – Grid-Forming Inverters and Hybrid Power Systems
   https://docs.nrel.gov
7. Uptime Institute – Data Center Capacity, Density, and Operational Trends
   https://uptimeinstitute.com
8. Power Stabilization for AI Training Datacenters
   Microsoft – Open AI - NVIDIA
9. AFCOM's 2026 State of the Data Center Report
   https://afcom.com/news/720658/The-Data-Center-Density-Dilemma.htm

6.2 Nomenclature / Abbreviations

Tech/Software	
IT	Information Technology
AI	Artificial Intelligence
ML	Machine Learning
EPC	Engineering-Procurement-Construction
BTM	Behind-The-Meter
TTP	Time to Power
BESS	Battery Energy Storage System
EMS	Energy Management System
PUE	Power Utilization Effectiveness
WUE	Water Utilization Effectiveness
GPU	Graphics Processing Unit
HVAC	Heating, Ventilation & Air Conditioning
LNG	Liquid Natural Gas
FOTM	Front-Of-The-Meter
GPUaaS	Gpu as a Service
Aero	Aeroderivative
GT	Gas Turbine
EMT	Electromagnetic Transient
EGT	Enhanced Gas Turbine
PV	Photovoltaic
GHI	Global Horizontal Irradiance
POI	Point Of Interconnection
WTG	Wind Turbine Generator
IBR	Inverter-Based Resources

Units	
kW	Kilowatt
MW	Megawatt
Sq-ft	Square-foot

ABOUT THE AUTHORS

Ihab Chaaban, P.Eng., MBA
Aeroderivatives Global Commercialization Director,
Data Centers Grid Firming, Hydrogen & Hybrids - GE Vernova

Ihab Chaaban is a global technical and commercial leader with over 30 years of experience in the Power Generation and Systems industry, 18 of which with GE, with deep domain expertise in the Energy, Utility and Electrical Power Businesses. Chaaban started his career as an Electrical Engineer and held several roles in Engineering and Services via global assignments in the Middle East, Africa, Germany, Brazil, Canada, and the USA.

Chaaban started his career with GE as an Application Engineering Lead for the Aeroderivative global division followed by other roles in commercial operations leadership and proposal management for Aeroderivative, distributed power and industrial frames gas turbines. Chaaban assumed his current role in GE Vernova to support the grid firming and data centers segments including Aero gas turbines in hybrid and green hydrogen applications with renewable assets and battery energy storage systems. Chaaban is a Licensed Professional Engineer in Ontario Canada, holds a B.Sc. in Electrical Engineering from Alexandria University, an MBA from Lansbridge University and speaks five languages (English, French, Arabic, Portuguese proficiently with a good command of Spanish).

Kassy Hart, M.Sc.
Aeroderivative Data Centers &
Hydrogen Commercialization Leader - GE Vernova

Kassy Hart is a technology strategist and engineer with 15 years of experience in the energy industry. She is responsible for the development of power generation opportunities for the Aeroderivative business in growth market segments, including global customer engagement, preliminary power plant configuration, commercial strategy and market & policy intelligence.

Kassy joined GE as part of the Edison Engineering Development Program. She spent her first 10 years in heavy-duty Gas Power Engineering with a focus on emergent technologies (additive manufacturing & repair; hydrogen fuel systems design) before serving in her most recent past position as Director of Energy Transition Technology, where her focus was on innovation strategy for GE Vernova HQ. Kassy holds a B.S. and M.S. in Mechanical Engineering from the University of Kentucky and completed the executive Global Energy Leadership Program at Rice University Business School. She is fluent in English with proficiency in Spanish, and holds 17 US patents.

ABOUT THE AUTHORS

Andrea De Simone, M. Sc.
Aeroderivative Data Centers &
Hydrogen Commercialization Leader – GE Vernova

Andrea De Simone is a distinguished global technical and commercial leader with more than 27 years of leadership experience in the power generation and energy sector across GE and its affiliates. In is current role, he is focused on advancing strategic solutions for the fast-growing data center segment, while also supporting commercialization initiatives in hydrogen.

His career spans engineering, innovation, program leadership, and business development, with deep expertise in gas turbines, power plants, energy systems, and utilities. Having held leadership assignments in Italy, Switzerland, and the United States, De Simone brings an international perspective to technology advancement and market development.

He began his career as an aerodynamic engineer specializing in compressors and turbines and went on to hold senior technical and managerial roles across industrial and heavy-duty gas turbines, power plant solutions, and innovation programs within the Gas Power sector. Throughout his professional journey, he has helped drive the integration of innovative technologies into power generation and supported strategic growth opportunities across an evolving energy landscape.

De Simone holds a Master of Science in Aeronautical Engineering from Sapienza University of Rome. He is fluent in English, Italian, and German, and has additional proficiency in Spanish and French.

GE VERNOVA

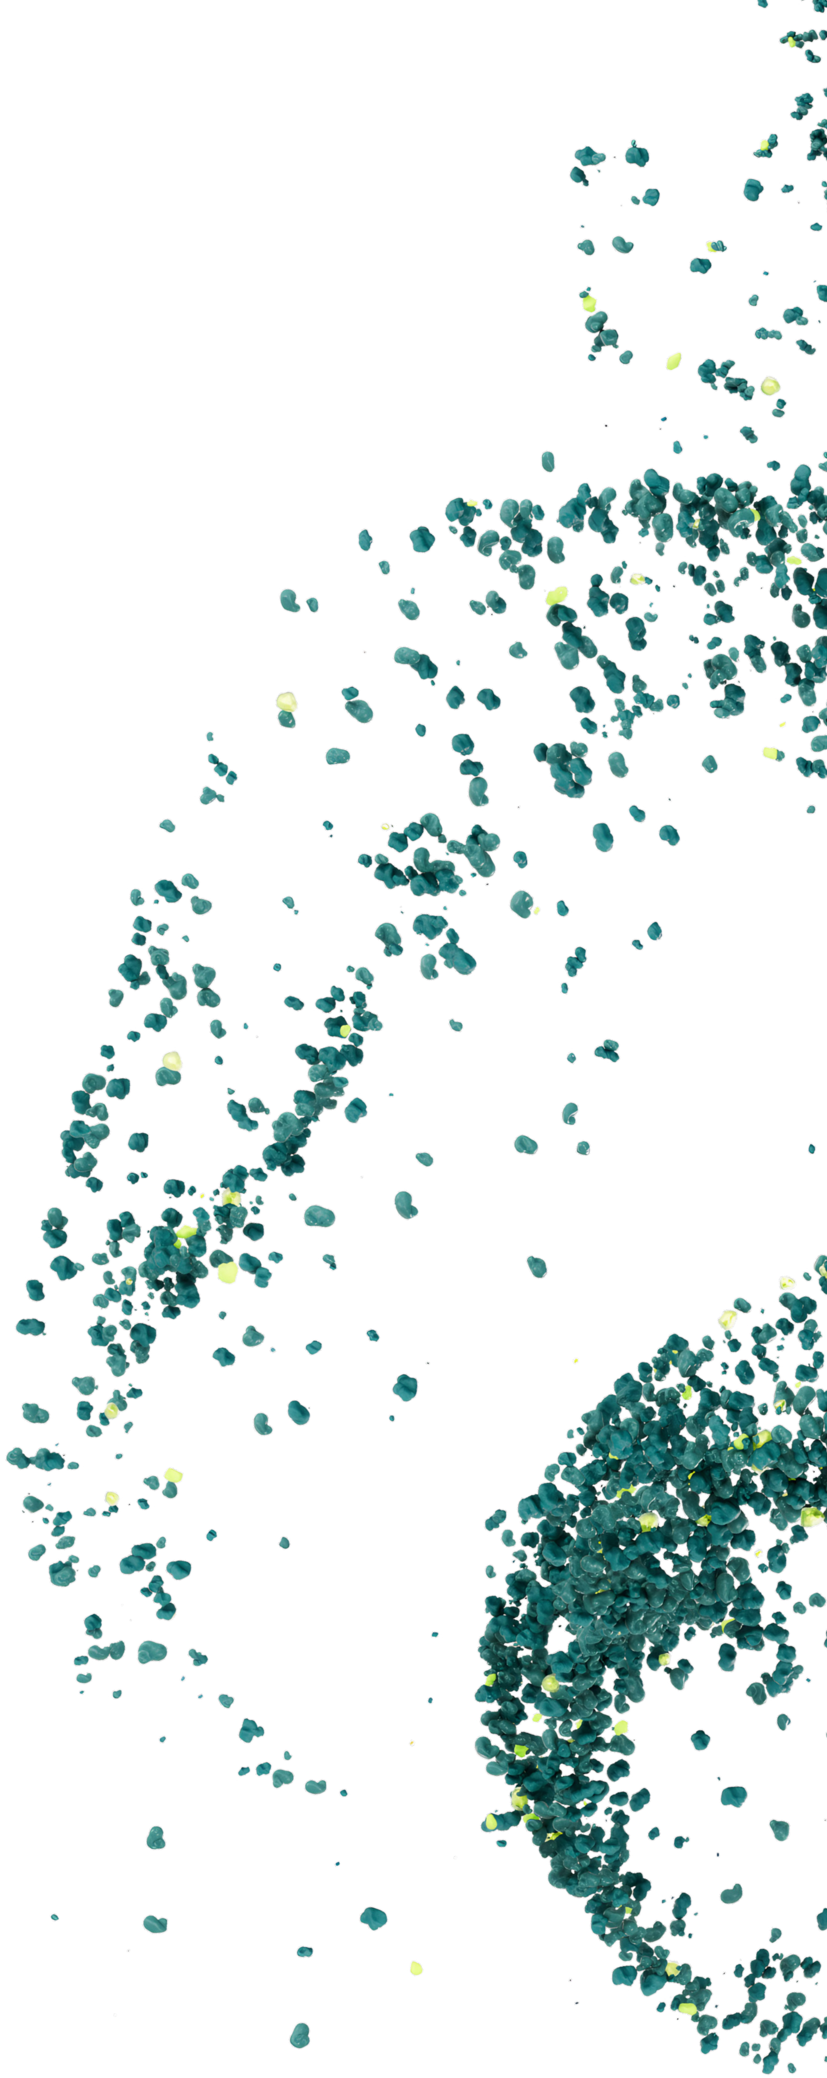

gevernova.com

GEA35974 07/2026