

From Backup Power to Grid Enabler

Use Cases for Aluminum-Air Generator (AAG™)
in AI-Era Data Centers

Clean Backup
Power for AI
Inference and
Cloud

Page 09

Unlocking Additional
Grid Capacity in
Power-Constrained
Markets

Page 11

An On-Demand
Resilience Layer
for Microgrids

Page 14

How AAG
Compares to
Alternative
Solutions

Page 15

Table of Contents

4	Executive Summary
5	The Changing Role of Onsite Reserve Power for Data Centers
7	The Aluminum–Air Generator <i>A clean and Reliable Onsite Reserve Power Source</i>
9	<i>Use Case 01</i> Clean Backup Power for AI Inference and Cloud Data Centers
11	<i>Use Case 02</i> Unlocking Additional Grid Capacity in Power–Constrained Markets
14	<i>Use Case 03</i> An On–demand Resilience Layer for Renewable Data Center Microgrids
15	How AAG Compares to Alternative Solutions
17	Conclusion



Audience & Scope

This white paper is intended for data center operators, power and infrastructure architects, and sustainability and technology leaders evaluating power resilience alternatives and architectures for their facilities. It is also relevant for utilities, regulators, and AHJs seeking to understand

next-generation backup technologies and their integration into data center environments. Business decision makers, engineering teams and system integrators focused on resilient, low carbon digital infrastructure will find value in the use cases and technical perspectives presented.

From Product to Deployment



The Phinergy-Rosendin collaboration

Phinergy is collaborating with strategic partners such as Rosendin Electric to roll out AAG solutions, combining aluminum-air generation technology with a proven electrical EPC delivery platform. Leveraging Rosendin's experience delivering complex, MW to GW, mission-critical systems, the collaboration provides a clear path from

pilot deployments to fleet-scale rollout, reducing adoption risk for data center operators. The partnership addresses key limitations of diesel backup systems, including emissions, permitting complexity, fuel logistics, and long lead times, while supporting the reliability, scalability, and schedule requirements of hyperscale and AI-driven facilities.

Executive Summary

Data center power architecture is entering a new phase. AI-driven growth is increasing the scale, and strategic importance of compute infrastructure, while grid interconnection timelines, local permitting, equipment supply chains, and community expectations are becoming material barriers to deployment. As AI-focused data centers expand and inference workloads become a larger share of demand, access to reliable, clean, and quickly deployable power is becoming one of the defining constraints on data center growth.

Beyond reliability and sustainability, speed-to-power is emerging as a strategic competitive advantage. A growing body of research indicates that limited demand flexibility during rare peak-demand periods can unlock disproportionately large amounts of grid connection capacity, allowing operators to energize facilities sooner while reducing the need for immediate network upgrades.

In this context, onsite long-duration generation is evolving from a backup asset into an infrastructure enabler. It is becoming a strategic infrastructure layer that can accelerate speed-to-power, support flexible grid interconnection, protect high-value AI workloads, and enable clean, community-compatible resilience.

Phinergy's Aluminum-Air Generator (AAG™) is positioned for this new role. By using aluminum as clean fuel, the AAG provides generator-class onsite reserve power with zero emissions, multi-day duration, long standby life, low maintenance, and modular scalability.

The most relevant applications for the AAG are the long-duration layers of the data center power stack: backup power, peak-capacity support, and microgrid resilience, where assets must remain idle for long periods and then deliver meaningful energy for hours or days without local emissions, noise, or fuel-management complexity.

This paper focuses on three high-value use cases where these characteristics are especially relevant: clean backup power for AI inference and cloud data centers; unlocking additional grid capacity in power-constrained markets; and an on-demand resilience layer for renewable data center microgrids.


The Changing Role of Onsite Reserve Power for Data Centers

Historically, backup power in data centers was designed around a relatively narrow operating model: short-duration UPS batteries bridge the transition from grid power to diesel generators, and diesel generators provide extended runtime during utility outages.

This architecture was effective because most data centers operated as relatively predictable loads, grid interconnection timelines were manageable, and emergency generator emissions were treated as an acceptable cost of reliability. That model is now under pressure.

Speed-to-power has become a strategic constraint

U.S. data center electricity demand is expected to grow sharply by 2030, with some states facing particularly high concentration of data center load.



Values are shown for 2024, 2030, and 2035 across three data center scenarios. Outputs from EPRI's US-REGEN model under reference policies. The figure highlights how rapid DC growth can raise the DC share from today's 4–5% to 10–20% by 2035. MDHD = medium- and heavy-duty vehicles; LDV = light-duty vehicles.

Source: EPRI – Powering Intelligence 2026

Traditionally, power availability was rarely the primary factor determining where and when a data center could be built. Today, access to electrical capacity has become a dominant constraint on AI infrastructure deployment. In many major markets, utilities cannot provide the full capacity requested by hyperscale developers because of generation, transmission, substation, transformer, or interconnection limitations.

The challenge is compounded by uncertainty surrounding future AI electricity demand. Utilities face a difficult dilemma: build infrastructure aggressively and risk stranded assets if projections prove too high or expand cautiously and risk slowing AI deployment and economic growth. Traditional generation and grid expansion may also require many years, leaving a widening gap between data center development schedules and utility delivery timelines.

Flexibility is changing the game

As a result, onsite reserve power assets are being asked to do more. They must still protect uptime, but they may also support flexible interconnection agreements, help bridge utility capacity shortfalls, complement renewables and batteries, and reduce the local environmental burden of dense data center clusters. This creates a market opening for long-duration, low-emission, low-maintenance backup technologies that can perform generator-like functions without the drawbacks of diesel combustion.



The Aluminum-Air Generator

A Clean and Reliable Onsite Reserve Power Source

An Aluminum-Air Generator (AAG) is an on-demand zero emission generation system that converts aluminum into electricity through a reaction with oxygen drawn from the surrounding air.


As with conventional generators, the aluminum inside the system is gradually consumed as energy is produced, and when it is depleted, fresh aluminum is loaded into the system to refuel it.

At the core of the AAG is Phinergy's aluminum-air technology, which includes aluminum-air cells with air electrodes serving as cathodes and aluminum plates serving as anodes, an electrolyte circulation system, a thermal management system, and an energy management system.

For data centers, this operating principle translates into an onsite reserve power system that behaves like a generator while avoiding combustion-based limitations. Because the AAG remains inert during standby and activates only when electrolyte is circulated through the cells, it can remain available for extended periods without self-discharge, fuel degradation, or routine fuel maintenance. Once activated, the same aluminum-air process provides long-duration, zero-emission power, making the AAG suitable for data center applications that require reliable multi-day resilience, low maintenance, and deployment compatibility with sensitive urban, grid-constrained, or low-carbon environments.

For a deeper explanation of the AAG technology, system architecture, aluminum-as-fuel concept, and deployment considerations, readers are encouraged to refer to the first white paper in this series: [Aluminum-Air Generator \(AAG™\) for Data Centers: Zero-Emission Long-Duration Backup Power](#)





Aluminum Air System Layout

Electrolyte is used to activate and deactivate the cells and for thermal management.



Use Case 01

Clean Backup Power for AI Inference and Cloud Data Centers


AI infrastructure is moving closer to end users

The first generation of AI infrastructure was largely focused on model training. These workloads were typically located in large remote campuses where access to land and power were prioritized over proximity to end users. By nature, these workloads were inherently more tolerant of outages, while their remote locations also made emissions a less significant deployment concern.

The current phase of AI growth is driven by inference, the continuous execution of trained models supporting applications such as enterprise AI assistants, real-time analytics, autonomous systems, healthcare platforms, financial services and industrial control systems.

Inference workloads require lower latency and are increasingly being deployed closer to the communities they serve. As data centers expand in scale and move closer to urban and suburban populations, the requirements for backup power begin to change.

According to McKinsey & Company, by 2030, 70–80% of AI demand is expected to be driven by inference, and related analysis describes a shift from predictable data center power usage toward highly volatile and power-intensive AI demand profiles. As inference becomes more latency-sensitive and geographically distributed, more compute capacity is expected to be deployed closer to demand centers.



The emerging conflict

Communities increasingly expect critical infrastructure to minimize air pollution, noise, visual impact, and environmental risk. At the same time, operators must maintain increasingly stringent uptime requirements. This creates a growing tension between reliability needs and community expectations.

Diesel generators remain widely used, but operators frequently encounter challenges related to air permitting, environmental reviews, community concerns, emissions restrictions, and noise limitations. These challenges can delay projects and increase development risk.

Diesel backup systems are increasingly visible in these debates because data center clusters may require hundreds or thousands of standby engines. In Virginia, for example, public reporting has

highlighted more than 10,000 backup generators serving data centers, with policymakers debating stricter requirements for cleaner generator tiers. Community concerns typically focus on local NO_x and particulate emissions, noise during testing or emergencies, cumulative air-quality impacts, and the perceived mismatch between clean-energy commitments and diesel-based resilience.

As AI services become embedded in critical operations, downtime can disrupt customer applications, enterprise processes, industrial systems, and public services. Consequently, operators require backup systems that deliver traditional generator-class reliability while remaining compatible with increasingly restrictive deployment environments.

The AAG - a clean backup power solution for the AI Inference era

The AAG was designed to provide long-duration backup capability while eliminating many of the characteristics that create challenges for community deployment.

The future data center must satisfy multiple stakeholders simultaneously: customers requiring uninterrupted service, utilities managing grid stability, communities seeking environmental protection, and governments pursuing decarbonization targets.

The AAG supports all four objectives within a single backup architecture.

Zero Local Emissions

The absence of NO_x, SO_x, and CO₂ emissions simplifies deployment and aligns with sustainability objectives.

Quiet Operation

Low noise and vibration make the technology suitable for sites located near residential and commercial developments.

Long-Duration Resilience

Unlike conventional batteries that are optimized for shorter discharge periods, the AAG provides multi-day backup capability.

Unlimited Standby

The system can remain dormant for extended periods without energy loss or fuel degradation.

Limited Footprint

The modular architecture enables deployment in environments where space efficiency is critical.

Key Takeaways

As AI compute moves closer to people, backup power must become cleaner, quieter, and easier to deploy. Long-duration resilience can no longer come at the expense of community acceptance. With clean, quiet, long-duration power in a compact footprint, the AAG directly addresses these requirements.

Unlocking Additional Grid Capacity in Power-Constrained Markets

Data center growth is increasingly limited by grid availability

Across many regions, access to power has become the primary constraint on data center development. Utilities are often unable to provide the full capacity requested by operators, not because annual energy is insufficient, but because the grid needs to cope with peak demand periods, contingency events, and transmission disruptions.

The scale of the constraint is material. According to EPRI, on an annual energy basis, U.S. data center electricity consumption could reach 380–790 TWh by 2030, representing 9–17% of total U.S. electricity

consumption. Data center load is also highly concentrated geographically: Virginia already represents an outlier, with data centers accounting for more than 25% of electricity demand, while several other states could cross the 20% threshold by 2030.

This concentration creates localized transmission, distribution, substation, and transformer constraints even where regional energy supply is sufficient on an annual basis.




Peak constraints often limit growth

Recent academic and industry research suggests that grid constraints are frequently driven by a relatively small number of annual peak-demand hours rather than by annual energy availability. Grid infrastructure must be planned for extreme conditions even when they occur only a few dozen hours per year. The IEA estimates that if data centers can reduce their grid consumption for only 0.1% to 1% of the time, existing power systems could

integrate substantial new data center capacity. In the United States, it estimates that up to 70 GW of grid capacity could be unlocked for data centers growth if they reduce their grid demand for 1% of the time. In Europe, it estimates that up to 40 GW could be similarly unlocked. This reduction can be achieved either by curtailing consumption, which may affect the quality of service, or by switching to onsite generation.

Data center capacity additions to 2035 and feasible integration into the current electricity system under different flexibility cases



Current electricity systems can already integrate all data center additions to 2035 if a mix of backup activation and workload management reduces grid demand 1% of the time

Source: Energy and AI, IEA, 2025

Targeted flexibility, major grid benefits

These findings are further supported by the Nicholas Institute at Duke University, which examined how limited, contractually controlled reductions in customer demand during periods of system stress could allow additional load to connect to the existing grid. Across 22 U.S. balancing authorities representing approximately 95% of the U.S. grid, its 2025 study estimated that nearly 100 GW of additional large loads could potentially be integrated using existing capacity with a 0.5% average annual curtailment rate and short-duration reductions in consumption.

Even limited flexibility can deliver major system benefits. Duke University's 2026 analysis found that avoiding grid demand during just 1% of peak hours – about 88 hours per year – could reduce new natural-gas capacity investment by approximately 10% to 50%, generating estimated savings of \$40 billion to \$150 billion over a decade.

The practical implication is precise: the value comes from releasing the exact constrained hours at the exact constrained location, not merely from reducing annual electricity consumption.

Onsite power generation as grid enabling asset

In power-constrained markets, long-duration backup systems can serve as contractual flexibility assets, temporarily reducing grid demand while maintaining onsite operations during periods of congestion, shortages, transmission constraints, or extreme weather.

Combining partial grid access with supplemental generation can support earlier interconnection and defer costly network upgrades, while preserving the same asset's emergency-backup function.

This approach transforms backup generation from an emergency asset into a capacity-enabling resource.

By combining operational flexibility with reliable onsite generation, a data center can evolve from a rigid source of power demand into a flexible grid asset, shifting from being part of the grid-capacity challenge to becoming part of the solution.

Key Takeaways

As power availability becomes the defining constraint on AI infrastructure growth, onsite reserve power can evolve into strategic tools for maximizing utilization of existing grid infrastructure and enabling faster data centers deployment. The AAG becomes a speed-to-power means by turning data centers into grid flexible assets.

The AAG - a speed-to-power enabler

Many technologies can provide short-duration support, but few are technically and economically optimized for supplying significant power for tens or hundreds of hours annually.

The AAG combines several characteristics that make it particularly attractive for this role:

Multi-day energy generation

Zero onsite emissions

Unlimited standby duration without performance degradation

Minimal maintenance requirements

Scalable deployment from individual facilities to large campuses

Unlike batteries, which become increasingly expensive as duration requirements grow, the AAG can economically provide large amounts of energy over extended periods.

Compared with diesel generators, the AAG avoids emissions, permitting complexity, and community opposition.



An On-Demand Resilience Layer for Renewable Data Center Microgrids

Renewable microgrids need a rarely used but critical onsite reserve power generation resource

Many next-generation data centers are pursuing highly sustainable power architectures built around solar generation, wind generation, battery energy storage systems, and advanced energy management platforms. These systems can dramatically reduce dependence on fossil fuels while supporting long-term decarbonization goals.

One of the major challenges in developing these microgrids is posed by the intermittency of renewable generation sources. These include seasonal changes (e.g. wind patterns), extreme weather conditions or unpredictable events (e.g. multi-day storms). A good renewable microgrid design balances the generation capacity with enough storage capacity to support the consumer throughout the average cases and with optimized capital expenditures. However, for the tail-end cases of the distribution curve, the renewable generation falls below demand for extended durations that can't be covered by storage capacity. Although these situations rarely occur, operators must still maintain supply to data centers operators.

The economic dilemma of tail-end reliability

The tail-end use cases require onsite on-demand generation capacity. These final few hours or days of reliability are often the most difficult and expensive to achieve. Technologies optimized for long standby and multi-day operation can address the extreme tail of reliability requirements without forcing operators to oversize batteries or generation assets.

Deploying diesel generators for addressing the above cases introduces emissions and fuel-management requirements that conflict with sustainability objectives.

Microgrid operators require a new category of assets: a clean energy system that may operate only occasionally but remains continuously available.

Positioning AAG as an on-demand reserve energy layer

The AAG is particularly well suited to this role as a zero-emission, multi-day generator with excellent standby characteristics. As such, it fits in the operational model, where under normal conditions, renewable resources and batteries meet all operational requirements. During prolonged renewable deficits or rare weather events the AAG is activated.

Key Takeaways

Renewable microgrids require on-demand generation to achieve reliability while not oversizing storage and intermittent power assets. The AAG provides that distributed reserve energy layer while supporting sustainability objectives.



How AAG Compares to Alternative Solutions

Examining the critical use cases discussed above shows that the common requirement is not simply backup power, but an onsite multi-day generation solution that can support data center resilience, grid-constrained operation, and renewable microgrid reliability. This solution must be clean, cost-effective, suitable for long standby periods, scalable at the site level, and practical to deploy within data center footprint, permitting, safety, and operational constraints. The following section compares the AAG with other technologies that could play this role, including diesel and biodiesel generators, natural gas generation, hydrogen fuel cells or turbines, and Li-ion batteries sized for long duration.

Key technology trade-offs

Diesel generators remain the incumbent solution because they are familiar, compact, dispatchable, and supported by mature supply chains. However, they also create the main constraints this paper addresses: onsite NO_x and CO₂ emissions, noise, fuel storage, fuel degradation, maintenance, testing, and air-permitting complexity. Biodiesel or renewable diesel may reduce lifecycle carbon intensity, but they do not eliminate combustion, local emissions, or generator permitting issues. Natural gas generation can support larger onsite power strategies, but it also remains combustion-based and depends on gas infrastructure.

Li-ion batteries and BESS are essential for UPS, AI power smoothing, and short-duration storage for integration of renewables. However, when the

required duration extends from minutes or hours to multi-day backup, batteries become increasingly constrained by cost, footprint, augmentation, fire-safety, and lifecycle considerations. The AAG is therefore addressing different needs than batteries: batteries provide immediate response and transient support, while the AAG provides the long-duration energy layer.

Hydrogen fuel cells and hydrogen turbines can offer low-emission power at the point of use, but their deployment depends on hydrogen availability, cost, delivery, onsite storage, safety frameworks, and infrastructure maturity. Other long-duration storage technologies, such as flow batteries, compressed air, gravity-based, thermal, or mechanical systems, may be valuable for grid-scale applications but can be difficult to deploy at dense data center sites due to footprint, siting, project scale, geography, or integration requirements.

Against this comparison set, the AAG occupies a distinct position: it provides generator-class, long-duration backup without onsite combustion; stores energy in inert aluminum rather than degrading liquid fuel or high-pressure hydrogen; enables continuous energy refill as needed, remains in standby without self-discharge; and can be deployed modularly at the site level. This makes it most relevant for the use cases addressed in this paper: clean backup for AI inference facilities, grid flexibility support in power-constrained markets, and on-demand reserve power for renewable microgrids.

This is summarized in the comparison table below.

ALUMINUM-AIR VS. ALTERNATIVES

Aluminum-Air Outperforms Alternative Clean Backup Solutions on TCO, Footprint and Sustainability

Technology	Total Cost of ownership (TCO) USD/MW · 3MW 48H · 10 H/YR	Power footprint M2/MW (INCL 48H)	Local NOx + SOx emissions KG/YEAR	Scope 1 + 2 emissions TON CO ₂ /MW	Safety Risk Profile
Aluminum-air	850-900k	20-30	0	0	Not Flammable Not explosive
Diesel genset	1,200-1,300k	22	300-320	650-700	Flammable Explosive
Biodiesel genset	1,250-1,400k	22	300-320	150-200	Flammable Explosive
Natural gas turbine	1,250-1,350k	30-32	5-10	160-200	Flammable Explosive
Fuel cell	2,250-2,400k	186-242	0	0	Flammable Explosive
H ₂ turbine	3,050-3,200k	186-242	5-10	0	Flammable Explosive
Li-ion	11,100-11,300k	222-278	0	0	Flammable Explosive
Iron-air	2,000-2,200k	650	0	0	Presumed Safe

Notes

- Assuming diesel gensets and natural gas turbines run for an additional 25 h/year for testing vs. other back-up options
- Assuming the H₂ gas is stored in a 300 bar tank
- Technology readiness level of low-NOx turbines is still not fully mature
- At large-scale installed capacities: 500 GW for diesel, 30 GW for biodiesel, 2.5 TW for natural gas turbines, 1 GW for fuel cells, 1 GW for H₂ turbines, 10 GW for Aluminum-air, and 5 TWh of Li-ion batteries

- Footprint of the power back-up system and gas storage, without considering grid connection
- Local onsite emissions, not considering upstream emissions
- Includes embedded emissions and assumes green electricity is used in the manufacture of Al, production of H₂, and charging of Li-ion.

Source: DoE, NREL, third-party consultant | Schneider Electric WP14

The table shows that aluminum-air is best positioned for the clean onsite long-duration generation layer because it combines zero onsite emissions, competitive footprint, and lower projected TCO than alternative solutions.

Conclusion

This paper examined three critical data center use cases: clean backup for AI inference and cloud data centers, unlocking additional capacity in power-constrained markets, and providing the last-resort resilience layer for renewable data center microgrids. It then compared the potential technologies that could support these requirements and demonstrated that the AAG is the best-positioned solution for the clean onsite long-duration generation layer in the data center energy structure.

The three use cases presented in this paper show where the AAG is particularly well suited. First, for AI inference and community-proximate data centers, the AAG enables long-duration backup power without onsite combustion emissions, noise, or fuel-management complexity. Second, in power-constrained markets, the AAG can help operators unlock additional effective grid capacity by serving as a clean, long-duration resource during periods when the grid cannot deliver the full requested load. Third, in renewable data center microgrids, the AAG can act as the final resilience layer, providing rarely used but critical multi-day onsite on-demand power when renewable intermittent generation and batteries are insufficient.

The comparison with other backup power technologies reinforces this positioning. Batteries

remain essential for UPS, ride-through, and AI power smoothing, but become less efficient as the required duration extends to many hours or days. Diesel and gas solutions are mature, but face increasing pressure from emissions, permitting, community acceptance, and sustainability objectives. Hydrogen and other long-duration alternatives may play important roles in the future, but many still depend on developing fuel infrastructure, siting conditions, or technology maturity. The AAG occupies a distinct position by combining generator-class long-duration resilience, zero onsite emissions, inert standby, low maintenance, and modular scalability,

Taken together, these characteristics position the AAG not only as a cleaner replacement for diesel generators, but also as an enabling technology for accelerating data center growth. It can support faster deployment in constrained power markets, improve compatibility with community and regulatory expectations, and strengthen the resilience of low-carbon and renewable-powered facilities. As data centers become larger, more power-intensive, and more strategically important, backup power solutions will need to deliver reliability, sustainability, and deployability at the same time. The AAG is designed to address this combined requirement.



References

1. Electric Power Research Institute. [Powering Intelligence 2026: Updated Scenarios of U.S. Data Center Electricity Use and Power Strategies](#). EPRI, Palo Alto, CA. (2026).
2. International Energy Agency. [Energy and AI. World Energy Outlook Special Report](#). IEA, Paris. (2025).
3. McKinsey & Company. Energy Storage for Data Centers. (2026).
4. Norris, T. H., Profeta, T., Patino-Echeverri, D., & Cowie-Haskell, A. [Rethinking Load Growth: Assessing the Potential for Integration of Large Flexible Loads in U.S. Power Systems](#). Durham, NC: Nicholas Institute for Energy, Environment & Sustainability, Duke University. (2025).
5. Phinergy. [Aluminum-Air Generator \(AAG™\) for Data Centers: Zero-Emission Long-Duration Backup Power Technology and Deployment Perspectives](#). (2026).
6. Washington Post. [Data centers boom, and Virginians breathe exhaust from 10,000 diesel generators](#). (2026, May 28).

Authors and Contributors

This paper was prepared by Phinergy in collaboration with Rosendin Electric, combining aluminum-air energy technology expertise with deep experience in the design and delivery of mission-critical electrical infrastructure for large-scale data centers and industrial facilities.



phinergy.com