Engineering the ENERGY AIR GAP™

Fuel Assurance for Hyperscale AI Environments

Context: The Return to Physical Reliability

AI-scale data center campuses are now being planned, financed, and operated like city-scale infrastructure. As developers shift toward behind-the-meter generation to protect schedule and uptime, the question moves beyond “megawatts on paper” to a simpler test: what actually runs when the system is stressed?

Our answer is fuel assurance built as physical infrastructure. At Cashman Preload Cryogenics (CPC), we call this the ENERGY AIR GAP™. It is an engineered, inside-the-fence buffer of onsite LNG storage and vaporization (sendout). This is not another ‘just-in-time’ delivery solution. It is a deliberate shift from kinetic logistics to static inventory. When conditions get tight, reliability comes from what you already control inside the fence, not from whether roads are passable, terminals are operating, or upstream systems remain unconstrained. The Energy Air Gap allows a mission-critical campus to maintain generation output through the system-defining hours, even when pipeline deliverability tightens, operating limits bind, or force majeure conditions appear. In simple terms: the Energy Air Gap replaces storm-time logistics with static onsite energy inventory.

Important clarification: this is not an “air gap” in the cybersecurity sense. It is a physical energy buffer inside the fence. It is a deliberate fuel assurance layer that turns reliability from a contractual assumption into a measurable capability (i.e., inventory, sendout capacity, redundancy, and testability).

Batteries manage seconds-to-hours. Diesel can cover hours-to-days, but becomes a permitting, maintenance, and storm-logistics burden at hyperscale. LNG storage is the scalable solution for multi-day run-through, providing high energy density, controllable sendout, and most importantly, fuel already onsite when the region is under stress.

This is not an average-day energy problem. Most hours, the system works. The risk concentrates in a small number of peak-driven, system-defining hours and multi-day weather windows when gas deliverability and grid constraints tighten at the same time, exactly when uptime matters most.

This is fundamentally a peak problem, not a baseload problem. Installed capacity is often sufficient across most hours of the year. The failures occur during a small number of system-defining windows, when peak demand, extreme weather, and fuel deliverability constraints coincide. Fuel assurance is about surviving those hours and days, not optimizing the average one.

Historically, electric utilities treated resiliency as an engineered obligation. They built in reserve margin, peak resources, and in cold-weather markets, physical fuel assurance like LNG peak shaving. Over time, deregulation and procurement norms shifted incentives toward utilization and just-in-time efficiency. That can work in normal conditions. But stress events keep exposing the same constraint: installed capacity can exist while the fuel path to run it is constrained.

The rest of this paper lays out the practical hierarchy of fuel deliverability architectures and a simple diligence test for developers, hyperscalers, and investors: can your campus run through the event window without depending on real-time pipeline flexibility or storm-time deliveries?

1. The New Paradigm: Reliability in the Era of AI-Scale Loads

The rapid expansion of AI-scale workloads has shifted power strategy for hyperscale developers from securing capacity to proving performance under stress. Traditional metrics and utility-style assumptions are not sufficient when the load is large, continuous, and operationally inflexible.

“System-defining hours” are the relatively small number of windows each year when the grid, gas system, or both become constrained (ie, extreme weather, transmission contingencies, and fuel deliverability tightness). During these hours, reliability is determined less by nameplate generation and more by the ability to convert fuel into power without interruption.

This risk is increasingly structural, not hypothetical. The National Petroleum Council’s recent work on gas-electric coordination highlights a persistent mismatch between how gas services are designed and how the power sector, and now large, 24/7 loads, need fuel to perform during stress. The implication for data centers is straightforward: a power plant without a defensible fuel assurance plan can become a stranded reliability asset or installed capacity that cannot run precisely when it is most needed.

The Energy Air Gap concept is a practical response to that reality. It does not replace the grid or the gas system. It creates a physical buffer that allows a campus to ride through curtailments, operating restrictions, and force majeure windows with fuel already inside the fence.

2. Evaluating the Hierarchy of Fuel Deliverability Architectures

Building a large behind-the-meter generation plant is a reliability promise to tenants. The risk is rarely that engines won’t run; it’s that fuel deliverability tightens exactly when you need it most (extreme weather, upstream disruptions, pressure events, or operator restrictions).

Most teams start with “firm gas” to support financing. That is often necessary, but it is not the same as a hyperscaler-grade fuel guarantee. No pipeline contract is written as a “five-nines” performance warranty, and physical constraints can persist across multi-day weather windows.

Single Pipeline + Firm Transport:

Securing firm transport on one line is often a bankable baseline. It improves priority in many conditions and can reduce exposure to routine curtailments.

Where it breaks is simple: firm rights improve priority, but they do not create molecules or eliminate physical constraints. In severe regional stress, upstream disruptions and hydraulic limits can still tighten deliverability, even when contracts exist.

Bottom line: firm-on-one-pipe is a prerequisite, but it is not an Energy Air Gap.

Dual Pipeline Interconnects:

Adding a second pipeline connection can reduce exposure to a single interconnect outage. The common misconception is treating "two pipes" as true N-1 fuel assurance.

Many high-impact events are correlated (deep freezes, hurricanes, widespread upstream disruptions, compressor/power interdependencies, regional operating controls). Two feeds can share upstream failure domains, so the second path may not be available when the first is stressed.

A firm + interruptible structure is not N-1. If conditions are tight enough to compromise the firm path, the interruptible path is typically already curtailed. A firm + firm approach can improve contractual priority, but true N-1 performance can require substantial duplicative take-or-pay commitments, while still leaving exposure to common-mode force majeure risk.

Takeaway: two feeds are better than one, but they can be expensive and still externally dependent during the system-defining hours.

Firm Gas + Diesel Backup:

Diesel can be an effective short-duration contingency, but at hyperscale it turns reliability into a scale and operating burden.

A 300 MW campus typically implies on the order of 100+ standby engines before redundancy, with significant testing, synchronization, and O&M overhead. Large diesel fleets also face rising permitting and community scrutiny in major data center markets.

The critical issue is long-duration reality. Multi-day operation shifts reliability into fuel management, tank farms, safety, and logistics. Extended runtime can require sustained storm-time deliveries precisely when roads, drivers, and terminals are least reliable.

Bottom line: diesel is a contingency tool. It is not a clean, long-duration resiliency architecture for AI-scale uptime.

Scale check (order of magnitude): Diesel at AI scale

A 1 GW campus running on diesel for 72 hours can easily require ~4-6 million gallons of fuel (illustrative).

At ~8,000-9,000 gallons per tanker, that is roughly ~500-800 truckloads if resupply is required.

Takeaway: If the design depends on storm-time deliveries, logistics becomes a primary failure mode.

Firm Gas + Virtual Pipeline of CNG/LNG:

Similar to the diesel fuel logistical burden, "virtual pipelines" use trucks to deliver compressed or liquefied gas to a site as a substitute for permanent, high-capacity deliverability.

Where it can fit:

- Early commissioning, bridging fuel while permanent infrastructure is built, or topping up onsite storage inventory.
- Smaller loads where the truck count is manageable and the exposure window is short.

Where it breaks at AI scale:

- Throughput: At hundreds of megawatts, daily fuel demand translates into a high-frequency trucking requirement that becomes a capacity, traffic, and coordination problem.
- Correlated risk: The same weather that creates the emergency often disrupts roads, drivers, terminals, and dispatch, right when deliveries must be flawless.
- Operational complexity: Safety, routing, permitting, and dispatch become part of the reliability stack.

Bottom line: Trucking can be a valid refill method when paired with onsite storage. It is not a durable "must-perform" fuel path at AI scale for the system-defining hours.

Interruptible Gas + LNG Storage:

This model uses lower-cost interruptible pipeline service and deliberately shifts the reliability burden onto onsite stored fuel.

The mindset is important: the pipeline becomes an economic refill path, not a must-perform lifeline during the tight hours. High reliability is achieved by sizing storage for a full run-through of the design-case window so the facility can ignore curtailments and operating restrictions when it matters most.

Tradeoff: this typically requires larger storage and disciplined inventory planning. The goal is not replenishment during stress; the goal is run-through. Refill occurs in off-peak or off-season periods when molecules and logistics are abundant.

Takeaway: this can be an excellent solution as hyperscalers become more familiar with LNG fuel assurance and are willing to design for run-through rather than contractual firmness.

Firm Gas + LNG Storage (The Gold Standard):

This provides the definitive Energy Air Gap: it combines the bankable baseline of firm transport with a physical, inside-the-fence hedge that is auditable and independent of real-time pipeline conditions.

Why it performs under stress is straightforward: the fuel is onsite, the sendout path can be designed with redundancy, and runtime is measurable. Inventory, sendout capacity, and test runs can be documented, turning fuel assurance from an assumption into a defensible capability.

Most importantly, during the event window you are not depending on storm-time logistics or real-time pipeline flexibility to survive the critical hours or days; you are operating on controlled inventory inside the fence.

Takeaway: for hyperscaler-grade expectations, this is the most defensible architecture because it reduces external failure domains and makes fuel assurance measurable, testable, and financeable.

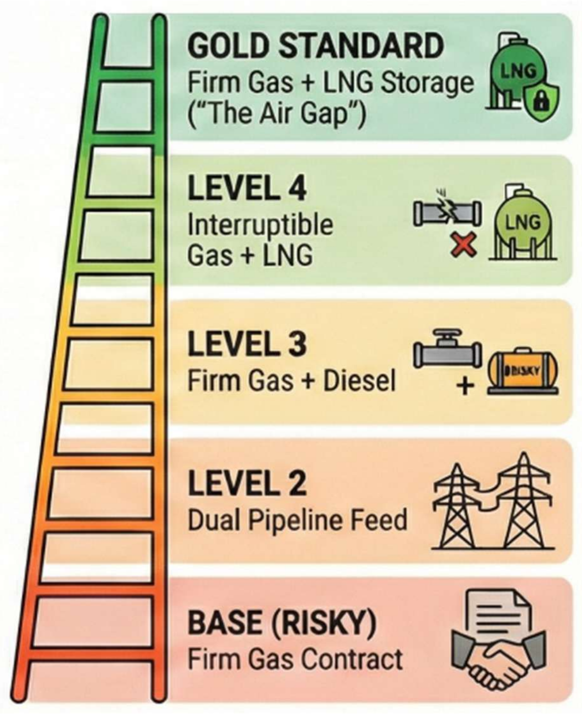

The Hierarchy of Reliability

Physical storage is the most direct way to reduce dependence on external constraints, turning fuel assurance from contractual hope into an engineered capability.

3. The Strategic Utility of Onsite LNG Storage

LNG storage is reliability infrastructure, not a commodity trade. Its primary value is making fuel assurance measurable, auditable, and defensible. By maintaining physical inventory and a controlled sendout path, a facility creates an engineered buffer that decouples runtime from real-time deliverability constraints.

The three (3) core functions of LNG fuel assurance are:

1. **Physical hedge:** Fuel is inside the fence, reducing dependence on external supply and pricing during constrained periods.
2. **Ride-through capability:** The site can maintain output for days (or longer, depending on storage volume) independent of real-time pipeline pressure and upstream disruptions.
3. **Audit-readiness:** Inventory, sendout capacity, redundancy, and test runs can be measured and documented for tenant diligence, lenders, insurers, and regulators.

A practical way to think about fuel assurance is duration. Batteries cover seconds-to-hours. Diesel can cover hours-to-days with significant operating and logistics complexity. LNG storage can be engineered for multi-day run-through with fuel already onsite.

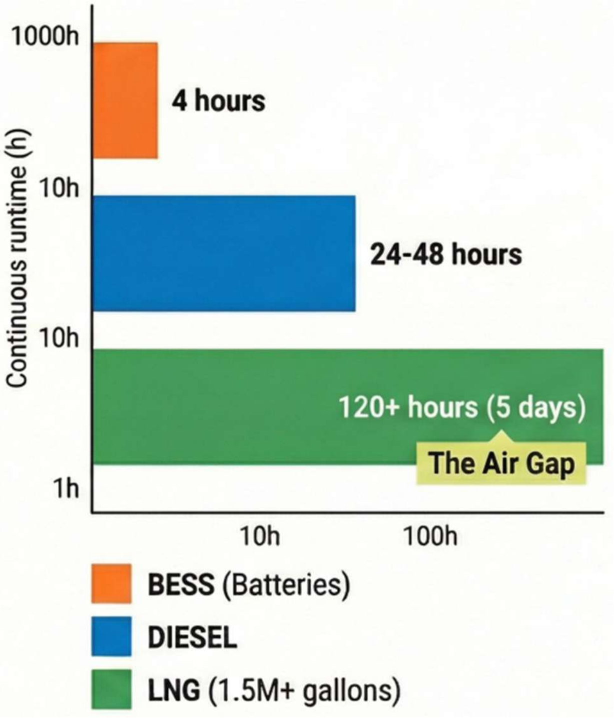

The Runtime Gap Chart

A critical component of this strategy is the "fill vs. refill" distinction. A fill strategy builds inventory ahead of tight seasons (winter and/or summer) so the site is independent of third-party logistics during the event window. Refill is an optimization performed when logistics and economics are favorable.

4. Deployment Logic and Practical Sizing

Sizing fuel assurance starts with one decision: what is your run-through requirement for the critical window, the duration you must be able to operate without relying on real-time pipeline performance or external logistics. The objective is simple: design inventory and sendout capacity so the facility can survive a system-defining event for a defined number of hours or days.

A disciplined sizing approach uses three (3) inputs:

1. Required electrical output (MW)
2. Technology class / heat rate (range-based early; refined later)
3. Ride-through duration (hours/days) you want to be independent during tight hours

Rule of thumb for early screening: 1,000,000 gallons of LNG ≈ ~70,000-75,000 MMBtu (gross), depending on composition and basis. Working inventory and heat rate drive "days of runtime," so final sizing requires a site-specific Basis of Design.

The right commercial question is not over-precision. It is:

"How many hours/days do we need to be independent of external gas constraints during the tight hours that matter?"

Table 2: Practical Guide to Sizing for Mission-Critical Fuel Assurance

CAMPUS LOAD (MW)	ILLUSTRATIVE GAS-BURN MMBtu per Day	24-HOUR BUFFER (MMGal of LNG)	5-DAY RUN-THROUGH (MMGal of LNG)	10-DAY RUN-THROUGH (MMGal of LNG)	20-DAY RUN-THROUGH (MMGal of LNG)
100	18,000	0.25	1.25	2.50	5.00
300	54,000	0.75	3.75	7.50	15.00
400	72,000	1.00	5.00	10.00	20.00
1,000	180,000	2.50	12.50	25.00	50.00
2,000	360,000	5.00	25.00	50.00	100.00

**Assumptions: nominal LNG basis ~74,000 Btu/gal (gross); ranges reflect heat-rate and operating-profile differences; final sizing requires a project-specific Basis of Design.*

Strategic note: In major stress events, fuel deliverability can tighten for multiple days. Contracts can include curtailment and force majeure provisions, and physical constraints can persist across a multi-day weather window. For many mission-critical designs, the “~5-day” run-through case is the pass/fail threshold that separates “we have gas most of the time” from “we can run through the event.”

Bottom line: define the run-through window first, then size storage so you can confidently treat pipeline performance during tight hours as a variable, not a dependency.

Deployment Models (how LNG fuel assurance is built):

Depending on siting, permitting, and the operating model, LNG fuel assurance is typically deployed in one of three patterns. Each is a different way to create an inside-the-fence buffer, with varying degrees of autonomy and system benefit.

Model A: Behind-the-Meter (BTM) Onsite Storage (“the air gap inside the fence”)

This is the most direct architecture. Storage and sendout sit on (or immediately adjacent to) the campus, supporting onsite generation with a defined runtime buffer. It is best suited for sites with the highest uptime requirements and the lowest tolerance for external dependency.

Model B: Regional Resiliency Hubs (“deliverability support at scale”)

In this model, LNG storage is developed as a strategic regional asset connected to the gas grid or positioned to supply multiple loads. The advantage is leverage: a single asset can improve resilience across several sites and support broader grid-and-gas reliability objectives.

Model C: Hybrid Models (“campus ride-through + upstream reinforcement”)

Hybrid designs combine onsite storage sized for immediate ride-through with regional storage or peak-shaving infrastructure that strengthens the upstream supply chain. This approach is particularly effective when hyperscalers, utilities/LDCs, and pipeline operators align on resiliency goals and operating protocols.

Takeaway: BTM storage maximizes autonomy, regional hubs maximize system leverage, and hybrid models often deliver the best combination (ie, campus-grade reliability with broader deliverability support).

A “fuel assurance” architecture can also be a “better-neighbor” architecture. When a campus can run through the tight window on controlled, onsite inventory, it is less likely to become an emergency participant in the broader system, chasing last-minute flexibility, depending on storm-time deliveries, or competing for constrained deliverability in the exact hours when communities are most exposed.

Optional (where coordinated and permitted): Some developers may choose to oversize storage and sendout beyond the campus’ minimum ride-through need so the asset can provide limited deliverability support to the upstream system during emergencies. That is not required for campus uptime, but it can convert a large load from a perceived risk into an infrastructure partner, especially in winter-sensitive times.

5. The Fuel Assurance Stress Test: A Due Diligence Framework

For hyperscale infrastructure, fuel assurance is no longer a checklist item. It is a pass/fail requirement for project viability. Teams should be able to defend their strategy against a simple stress test.

The 5-Day Run-Through Test (simple pass/fail)

If onsite generation is part of your uptime story, ask one question early:

Can we sustain design output for five consecutive days without relying on (a) storm-time deliveries or (b) real-time pipeline flexibility?

Why five days? The highest impact events are rarely one-hour problems; deliverability can tighten for multiple days.

If the answer is no: either accept derate/load shed during a regional event or engineer inside-the-fence fuel assurance sized for run-through.

- **Tight-hour performance:** What is the documented plan to maintain output during the most constrained hours of the year (extreme cold/heat, transmission contingencies, and fuel deliverability tightness)?
- **Ride-through duration:** Can the facility sustain design output for the duration of a plausible multi-day curtailment or force majeure window without relying on deliveries or real-time pipeline flexibility?
- **External failure domains:** Have upstream dependencies (supply chain, compressor/power interdependencies, operating restrictions, regional curtailments) been mapped, and reduced by inside-the-fence capability?

- **Auditability:** Are inventory, sendout capacity, redundancy, and test runs measurable and verifiable through physical testing, rather than only contractual language?
- **Operational dynamics (secondary, but real):** If the campus has meaningful ramps or step changes, is the fuel path engineered to avoid generator derates/trips once the event extends beyond the first seconds-to-minutes?

6. Conclusion

AI-scale data centers are no longer edge loads; they are city-scale infrastructure. In this environment, reliability is no longer determined by nameplate capacity or contractual priority. It is determined by whether fuel can be converted into power during the system-defining hours when both grid and gas systems are under stress.

The Energy Air Gap reframes fuel assurance as physical infrastructure. By combining onsite LNG storage with engineered sendout, it creates a measurable, testable buffer inside the fence. The Energy Air Gap reduces dependence on just-in-time pipeline flexibility and storm-time logistics when they matter most.

The diligence question is simple: can your campus, your city-scale load, run through a multi-day stress event without depending on real-time gas deliverability or emergency deliveries?

If the answer is no, then fuel assurance remains an assumption, not a capability.

At AI scale, reliability is no longer about contracts; it is about what you own, inside the fence.

About Cashman Preload Cryogenics (CPC)

Cashman Preload Cryogenics is a U.S. LNG storage tank and fuel assurance infrastructure company specializing in medium- to large-scale cryogenic storage systems for energy resiliency applications, including behind-the-meter generation supporting hyperscale data center campuses.

Publishing note: Prepared for discussion purposes only. Site-specific engineering, permitting, emissions compliance, safety analysis, and operating procedures are required.

Appendix A: Digital Speed vs. Hydraulic Physics

This appendix is included for teams that want the operational intuition behind one common failure mode. Digital loads can change quickly; pipelines are pressurized fluid systems that respond over longer time constants and within operating rules. When a large behind-the-meter plant ramps quickly, it can create a short-term mismatch between instantaneous fuel demand and the local system’s ability to replenish pressure at the same rate.

Batteries and controls typically cover the first seconds-to-minutes of electrical transients. The role of LNG sendout is different: it is not spinning reserve, but it can be engineered to come online within several minutes (i.e., ~5-10 minutes) depending on configuration to stabilize fuel delivery if local pressure, operating limits, or curtailments begin to bind.

The key point is not “daily balancing.” The key point is ride-through: LNG storage provides an inside-the-fence fuel layer that can carry the facility through the system-defining window when upstream deliverability tightens.

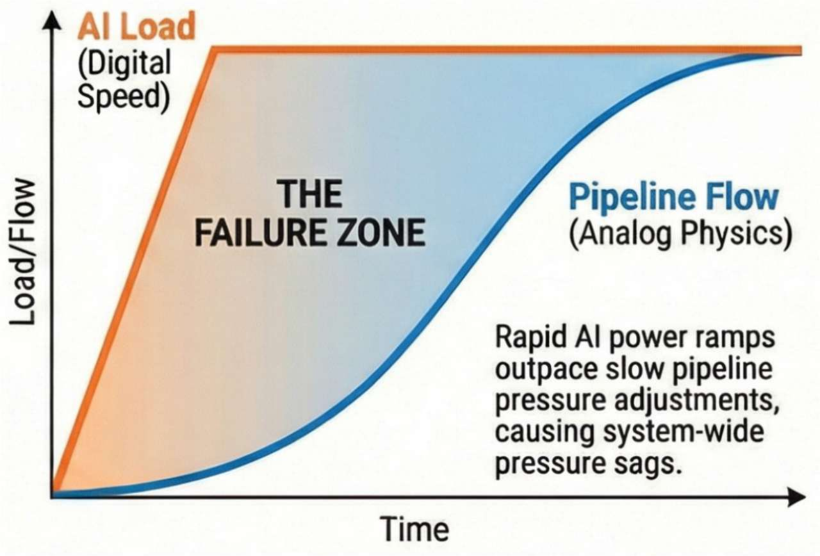

Table A1: The Failure Zone - Demand Ramps vs. Infrastructure Response

Component	Operational Speed/Characteristic	Failure Point Under Stress
AI Training Cluster	Seconds to Minutes (Step Function)	Sustained high utilization can exceed the instantaneous fuel arrival rate.
Natural Gas Pipeline	Hours/Days (Ratability & Scheduling)	Local pressure at the lateral/regulator station can sag under stress.

Onsite Gas Regulation	Mechanical/Automated Response	Plant can derate or trip on low fuel pressure if not buffered.
System Operations	Rules-based / Manual Interventions	Operational Flow Orders (OFOs) or hourly restrictions can limit flexibility.

The Volatility Shield

Illustrative market spike vs. pre-built inventory cost basis (site-specific)

The Energy Air Gap doesn't just protect uptime; it protects OpEx. You fill in summer; you ride through winter.

CASHMAN·PRELOAD
CRYOGENICS

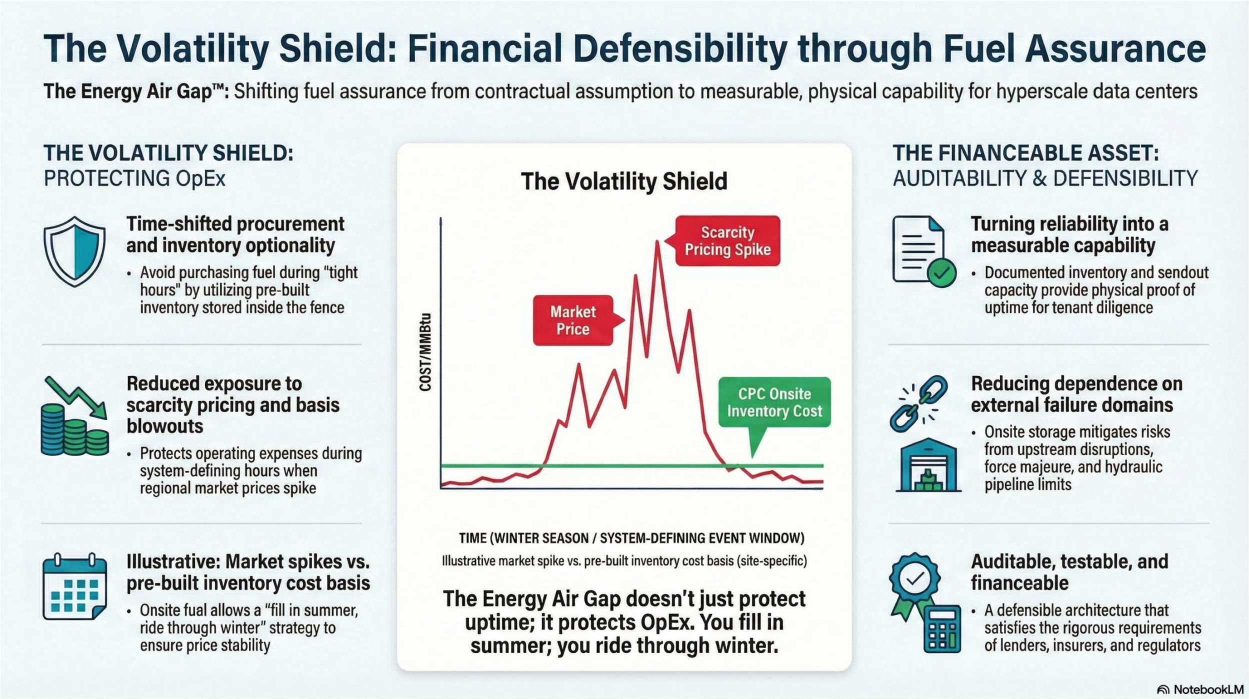

The Volatility Shield: Financial Defensibility through Fuel Assurance

The Energy Air Gap™: Shifting fuel assurance from contractual assumption to measurable, physical capability for hyperscale data centers

THE VOLATILITY SHIELD:
PROTECTING OpEx

Time-shifted procurement and inventory optionality

- Avoid purchasing fuel during "tight hours" by utilizing pre-built inventory stored inside the fence

Reduced exposure to scarcity pricing and basis blowouts

- Protects operating expenses during system-defining hours when regional market prices spike

Illustrative: Market spikes vs. pre-built inventory cost basis

- Onsite fual allows a "fill in summer, ride through winter" strategy to ensure price stability

The Volatility Shield

Illustrative market spike vs. pre-built inventory cost basis (site-specific)

The Energy Air Gap doesn't just protect uptime; it protects OpEx. You fill in summer; you ride through winter.

THE FINANCEABLE ASSET:
AUDITABILITY & DEFENSIBILITY

Turning reliability into a measurable capability

- Documented inventory and sendout capacity provide physical proof of uptime for tenant diligence

Reducing dependence on external failure domains

- Onsite storage mitigates risks from upstream disruptions, force majeure, and hydraulic pipeline limits

Auditable, testable, and financeable

- A defensible architecture that satisfies the rigorous requirements of lenders, insurers, and regulators

NotebookLM