VALIDATUS | VOLTAXIS

STRATEGIC DECISION INTELLIGENCE · ENERGY · DATA CENTRE INFRASTRUCTURE

Voltaxis Asked a Hardware Question. The Data Answered With a Contract.

Voltaxis is weighing roughly €150M over 24 months to enter the behind-the-meter storage market for AI data centres in the United States and Europe. The demand evidence is loud. But the analysis returns an inversion: the buyers driving this market prefer not to own power assets at all. What enters this market is a contract, not a battery, and a 90-day kill test decides whether anyone will sign it.

CURRENT DECISION POSTURE

CONDITIONAL GO

Hold the scale commitment. Fund the validation sprint.

The demand case needs no argument: data-centre electricity consumption is set to more than double by 2030, and the grid cannot keep pace. The finding that decides the €150M sits elsewhere. **The decisive competitor in this market is a contract structure, not a rival battery. Data-centre operators prefer not to own power assets, meeting new power needs as offtakers under purchase agreements, and the analysis ranks conventional hardware last among their real alternatives. Their preferred instrument is also failing them: the renewables PPA pipeline covers only around 15% of the demand growth now arriving.** What enters this market is therefore a service contract, and in procurement this sticky, the first signed, paid contract is not a sale. It is the position this market will judge every later sale against. The €150M is not the price of entry; it is the prize for winning that contract inside 90 days.

WHAT MUST BE TRUE BEFORE SCALE CAPITAL MOVES

01 **Value proposition validation.** A signed, paid pilot agreement with a target hyperscaler. *·90 days · Reversible*

02 **Business model viability.** The storage-as-a-service financial model validated by a target customer's finance and procurement team. *·6 months · Reversible*

03 **Whole product completion.** The pilot completed against the customer's own technical and operational acceptance criteria. *·9–12 months · Partially reversible*

This is a starting position, not a final verdict.

VALIDATUS ENERGY ANALYSIS · JUNE 2026 · VOLTAXIS DATACORE

5	50 min	143	111
DIMENSIONS	ANALYSIS TIME	EVIDENCE ITEMS	SOURCES

SECTION 01 · ABOUT THIS ANALYSIS

What this document is, and what it can settle

This report answers one question: should Voltaxis commit roughly €150M over 24 months to enter the behind-the-meter battery storage market for AI data centres, or remain focused on its existing utility-scale business? It was produced by the Validatus engine from one structured brief and the evidence base described below, then audited claim by claim against that evidence.

THE BRIEF — VERBATIM FROM THE ENGAGEMENT SCOPE

"Whether Voltaxis should invest about €150M over 24 months to enter this hyperscaler/data-centre BTM storage market, or remain focused on its existing utility-scale BESS business."

Product. Voltaxis DataCore, a proposed behind-the-meter battery energy storage system for AI data centres, with 8–12 hour duration and a bundled software layer for predictive load management and grid arbitrage, as described by the client. **Target.** Hyperscalers and large colocation operators running AI training and inference facilities. **Geography.** United States and Europe.

The evidence base. The analysis drew on two client-supplied documents, the Lawrence Berkeley National Laboratory 2024 United States Data Center Energy Usage Report and the International Energy Agency's Energy and AI special report, alongside 111 third-party sources. Across 28 analytical factors it harvested 143 evidence items, of which 26 factors produced computable findings.

Limitations, stated upfront. These are what make the verdict verifiable, not what weaken it.

LIMITATION 01

Two factors carried no computable evidence. One business-model factor within the Product dimension and the satisfaction-and-feedback factor within Experience returned no validated evidence, leaving those dimensions at nine-of-ten and two-of-three factor coverage respectively. Claims in those areas are correspondingly conservative.

LIMITATION 02

Thirteen unresolved evidence tensions. Across four of the five dimensions, the two evidence streams disagreed on a defensible single reading thirteen times, including on demand strength and entry timing, where the client's framing ran ahead of the public evidence. These tensions are penalised in the scores rather than smoothed over.

LIMITATION 03

A directional baseline, by design. Evidence confidence is moderate to low across all five dimensions, and no product documentation was part of the run: every product capability statement in this report is the client's own description. The evidence is also geographically uneven: the brief spans the United States and Europe, while the queue data, deal record and regulatory texture in this report are predominantly American. The structural read is firm; the precision of individual figures, and their transfer to the European half of the brief, is not.

This assessment is based on public and proxy inputs. Additional internal data would materially improve precision and relevance. The five dimensions point at the same strategic shape under different evidence: cross-dimensional agreement is high, and the overall analytical confidence reading for that layer is 0.92.

HOW TO READ THIS DOCUMENT

Page 3 is the verdict at a glance. Pages 4–5 name the three structural forces. Page 6 names the position they create. Pages 7–8 read the five scores. Pages 9–10 set the three gates. Pages 11–13 cover the buyers, the competitive frame, the scenario weights and the structural risks. Page 14 sequences the next 12 months. Pages 15–16 name what this baseline cannot answer; page 18 records the method and sources.

Confidence values throughout this document sit on a 0.0 to 1.0 scale: how strongly the evidence base supports the claim. They reflect data depth, not certainty.

EXECUTIVE SUMMARY

VERDICT · CONDITIONAL GO

THE ONE FINDING

Voltaxis asked a hardware question. The data answered with a business-model one: the decisive competitor here is a contract, not a battery — and the market is readier for the product than the seller is yet shaped to enter it.

The buyer prefers not to own power assets, so the strongest competitor is the purchase agreement and the binding constraint is the absence, today, of a signed contract in Voltaxis's name. **The €150M is not the price of entry; it is the prize for the first signed pilot.** The two tables below are the evidence; the rest of the document is the architecture.

EVIDENCE I · THE FIVE SCORES

base → final; two movements carry the finding

Brand

▲ RISES +15

Base	30
Final	45

The positioning whitespace is real and the strategy fits it. An empty niche, rewarded.

Market

▼ FALLS −16

Base	45
Final	29

Real demand, but evidence contested in seven places the scoring penalises rather than smooths.

Product

33 → **34** · one factor strong, no reference yet

Consumer

39 → **28** · buyer-fit, willingness-to-pay open

Experience

28 → **24** · provisional, evidence thin

EVIDENCE II · THE TWO-YEAR SCENARIOS

two futures tied; one instruction

◆ RISK-AVERSE · DOMINANT

AGGRESSIVE GROWTH

DIFFERENTIATION

The tie
The two most-weighted futures are level at 39%. Probabilities are directional model output, not forecasts.

The instruction is identical
In a risk-averse market the paid pilot is the entry ticket; in a growth market it is the head start.

Robust to the tie
The gates hold under both; Section 09 develops both alongside the two structural risks.

The market is more ready than the seller. **Everything that follows closes that gap with proof a pilot creates** *— the forces explain the reframe, the position page names the slot, the scores show the arithmetic, and the gates name what must be signed before the €150M moves.*

SECTION 02 · THREE FORCES SHAPING THE DECISION

The demand is not a forecast anymore. It is a queue.

FORCE 01 · TAILWIND · IRREVERSIBLE 2–3 YEARS

AI has turned data-centre power from a procurement line into a constraint on growth.

Global data-centre electricity consumption is set to more than double, to around **945 TWh by 2030** from roughly 415 TWh in 2024, in the IEA's base case. That end-state is slightly more than Japan's total electricity consumption today. The driver is the accelerated-server fleet behind AI workloads, whose electricity use the IEA projects to grow 30% per year, against 9% for conventional servers.

The capital is already moving. Alphabet, Amazon, Meta and Microsoft were reported to be planning as much as **$300 billion of AI-related capital expenditure in 2025**, and Goldman Sachs Research projected in 2024 that data-centre power demand would grow 160% by 2030. In the United States, the LBNL national assessment measured data-centre consumption at 176 TWh in 2023, which was 4.4% of all US electricity. Its 2028 projection runs from 325 to 580 TWh, between 6.7% and 12.0% of the national total.

The grid cannot keep that pace. The IEA's analysis finds that grid constraints could delay around **20% of the global data-centre capacity planned for construction by 2030**, and by mid-2026 industry outlooks put grid-connection waits in primary data-centre markets **beyond four years**. Power availability, not land or chips, is becoming the binding constraint on AI build-outs.

WHAT THIS MEANS FOR THE DECISION

A product whose stated purpose is speed-to-power addresses the precise point where this force binds. The demand side of the Voltaxis DataCore thesis does not need to be argued. The open question is what form a seller must take for this buyer to act, and the next two forces answer it.

IEA, Energy and AI (World Energy Outlook special report) · LBNL, 2024 United States Data Center Energy Usage Report · Goldman Sachs Research, May 2024. Dollar figures as reported in the source documents (USD).

SECTION 02 · THREE FORCES — CONTINUED

FORCE 02 · STRUCTURAL FRICTION THE ORGANISING FINDING · ACTIVE NOW

The buyer does not buy what Voltaxis plans to sell.

"It is clear that owning power assets is not the preferred business model for data centre operators." The sentence is the IEA's, not this report's: these companies meet new power needs indirectly, as offtakers under corporate power purchase agreements, keeping generation and storage off their own balance sheets. The competitive analysis runs with the same grain. Ranked by value to the buyer, the real alternatives to DataCore are **the purchase agreement first, doing nothing second, grid expansion third, and conventional storage hardware last**. The strongest competitor in this market is a contract structure.

That structure is also cracking. The IEA finds that the renewables PPA pipeline, projects under development plus those announced, would cover only around **15% of projected data-centre electricity demand growth to 2030**. The buyer's preferred instrument is dominant and insufficient at the same time, and that gap is the opening.

> **WHAT THIS MEANS FOR THE DECISION**
>
> The product question and the business-model question are not two questions. DataCore enters as storage-as-a-service, priced in the buyer's commercial language, or it does not enter. The brief priced a hardware entry; the data prices a contract.

FORCE 03 · STRUCTURAL FRICTION 12–18 MONTHS TO WORK THROUGH

The wall around this market becomes the moat inside it.

Energy infrastructure design and operation traditionally prioritise reliability and safety over rapid innovation, and the evidence describes conservative, switching-cost-heavy procurement in which an unproven supplier faces a long sales cycle however urgent the need. Adjacent industrial behaviour shows how far that loyalty runs: in a 2019 survey of 146 large manufacturing and process firms, **71% renewed their existing energy-management provider rather than switch**. Read the same stickiness from the other side and it changes meaning. Once a service contract is signed, the conservatism that walls Voltaxis out today can start working for it. One calibration the proxies cannot supply: the loyalty data comes from industrial buyers, and the largest hyperscalers deliberately multi-source to avoid lock-in, so the moat reading is strongest one tier down, with colocation operators and embedded service relationships. At the very top of the pyramid the first contract is less a lock than a key: the reference this market demands before it talks to anyone. That is why the deepest condition the analysis sets for scaling is validation of a retention moat, and why the game it names is domination of a narrow niche rather than broad share capture.

> **WHAT THIS MEANS FOR THE DECISION**
>
> Gate 1 is not only validation. It is a race to be inside the fence while the category is still forming, because in this market the first paid contract is the position.

Read together, the three forces define the game. The demand force says move now. The contract finding says move in the buyer's commercial language or not at all. The stickiness finding says whoever is first inside an account is very hard to remove, in either direction. This is a validation game whose prize is occupancy.

IEA, Energy and AI · Energy Policy (2019 survey, 146 firms) · Engine cross-dimensional synthesis.

SECTION 03 · THE POSITION

The contracted firming slot, and why it is still empty

The forces in Section 02 create an opening; this page names where it sits. The buyer's real power options arrange themselves on two axes: whether the asset touches the buyer's balance sheet, and whether the power sits on the buyer's site. Every option this market currently uses concedes one axis to win the other.

THE OPTION	WHAT IT OFFERS	WHAT IT LEAVES OPEN
Corporate power purchase agreement	A contract, off the balance sheet	Off-site: it answers the emissions ledger, not the interconnection queue or the resilience requirement, and Section 02 shows its pipeline covers a fraction of the demand now arriving.
Grid connection plus conventional UPS	The default	Years in the queue for the connection; minutes of cover from the batteries.
On-site gas as primary power	Fast, proven, and deploying at scale in 2026	Sits on the wrong side of public clean-energy commitments, and cannot smooth volatile AI loads or provide the instantaneous response that flexible-interconnection deals reward.
Contracted on-site firming — the forming slot	On-site control, delivered and guaranteed as a service contract	No supplier has yet claimed the AI-native version of this position. **This is the slot DataCore is aimed at.**

THE WEDGE — THREE CONVERGING CONDITIONS

01 **The buyer prefers not to own power assets.** Whatever fills the slot must be a contract; the IEA states the preference plainly and the 2026 deal structures confirm it.

02 **The buyer's preferred contract is insufficient and placeless.** The PPA pipeline covers a fraction of coming demand growth and puts nothing on the buyer's site, so the contract has to move behind the meter to answer the actual constraint.

03 **The on-site boom needs exactly what gas cannot supply.** Smoothing for volatile AI loads, instantaneous response, clean firming, and the flexibility that unlocks earlier interconnection. The client-stated duration and predictive software layer are aimed at this role, and the firming slot is the defensible one rather than the primary-power fight. One open edge, taken up in Section 10: the first deals in this slot run on far shorter systems than the client specifies, so the pilot must price the duration, not assume it.

THE POSITION

The position is the answer to a question this market is now asking out loud: who will put firm, clean, AI-shaped power on my site without putting it on my balance sheet, and sign for its performance? Through the three gates, the answer Voltaxis can become is: we do.

SECTION 04 · WHAT THE ANALYSIS SCORED

The 29% market score is not a verdict

The scores price the distance between that position and the company currently standing outside it. Two movements in this table look wrong until the method is visible, and both carry the strategic finding. The Market dimension enters with the strongest base reading, 45, and leaves at 29, with a caution flag. The penalty is not weak demand. It is conflict: seven unresolved readings inside the market evidence, including places where the client's framing of demand strength and entry timing ran ahead of what the public evidence supports. The scoring prices disagreement, so a loud market with contested claims scores below a quiet one with clean ones.

The Brand dimension makes the opposite journey, entering near the bottom at 30 and finishing highest at 45. That is not because the brand is strong. It is because the brand strategy is: the positioning whitespace for a purpose-built, utility-grade entrant is real, and the strategy adjustment rewards a credible plan applied to an empty position. Both movements say the same thing. The market is not rejecting the product; it is pricing the absence of proof, and proof in this market has a precise shape: a signed service contract.

FIVE DIMENSIONS · BASE → FINAL SCORE

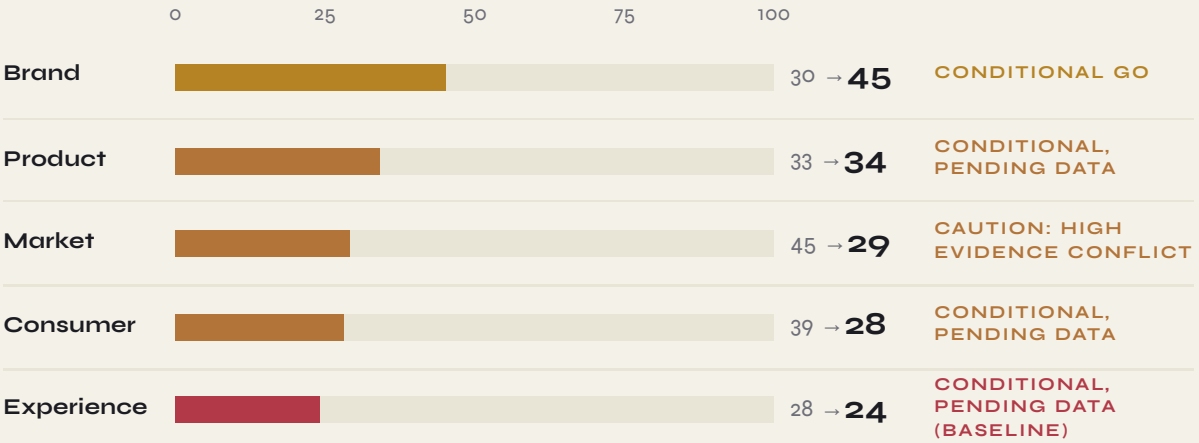

META READING 01

Is the market ready for this category?

53 MARKET READINESS

The demand side is prepared: the constraint is real, the buyers are concentrated, the pain is current.

THE 16-POINT GAP

Directional model readings

META READING 02

Is the seller ready for this market?

37 DIFFERENTIATION

One calibration of the same finding: the market is more ready than the seller, chiefly because the seller is not yet shaped like the thing this market buys.

A separate readiness tool, scoring the decision across six lenses, returns the same inversion in rank order: market timing highest, business-model viability lowest. Directional, but it is the spine arriving from a third direction.

Final scores apply a strategy-fit multiplier and a penalty for unresolved evidence tensions to each base score; the arithmetic for each dimension is shown on the following page. Tension counts on these two pages use the finest evidence-layer granularity in the analysis; the thirteen tensions in Section 01 are the same disagreements aggregated to factor level. Confidence values sit on a 0.0 to 1.0 scale and reflect data depth, not certainty.

SECTION 04 · FIVE DIMENSIONS, FIVE READS

Brand

30 → 45

Strategic read: build category credibility from scratch, borrowing the parent company's utility-grade reputation.

Score arithmetic: base 30% × 1.58 strategy fit − 2.0pp tension penalty = 45%.

LIFTED BY moderate positioning and differentiation strength in an empty niche; a named-competitor field with no specialist incumbent.

HELD BACK BY near-zero market presence, with cultural footprint very low and brand equity low; absence from the channel's analyst and trade coverage.

TO RAISE CONFIDENCE third-party validation, technical whitepapers, and one referenceable customer; the archetype evidence points the same way, an authority posture earned through proof rather than challenger noise.

Product

33 → 34

Strategic read: dominate a narrow niche, AI-native behind-the-meter storage, rather than contest the broad storage market.

Base 33% × 1.23 − 7.0pp = 34%.

LIFTED BY business-model resilience and durable-demand factor rated high, the strongest factor inside the Product dimension; client-stated 8–12 hour duration aimed at a real gap.

HELD BACK BY no reference deployment; seven unresolved evidence readings, entry timing among them; one business-model factor with no computable evidence.

TO RAISE CONFIDENCE an independent performance audit and a completed pilot against customer-set criteria.

Market

45 → 29

Strategic read: a direct contest with entrenched alternatives, grid power, purchase agreements and conventional backup, for the same budget line.

Base 45% × 0.79 − 7.0pp = 29%, caution flag for evidence conflict.

LIFTED BY demand and growth factors, both rated high, the verified demand surge of Section 02.

HELD BACK BY market stability and risk, rated low; seven unresolved readings, including demand-strength claims that exceed the public evidence.

TO RAISE CONFIDENCE pilot-derived demand data to replace contested projections.

Consumer

39 → 28

Strategic read: win operators away from incumbent power models, which means defeating the default option, not a rival vendor.

Base 39% × 0.78 − 2.0pp = 28%.

LIFTED BY moderate buyer-priority alignment: reliability and long-term cost dominate industrial purchase criteria.

HELD BACK BY operator distrust of AI in system-critical roles; willingness to pay unquantified; demand-signal tension between projected need and current adoption behaviour.

TO RAISE CONFIDENCE direct evidence from 5–10 operations directors on ranked fears and acceptable proof.

Experience

28 → 24

Strategic read: Provisional — baseline assessment. The analysis declined to assign a strategy on this evidence and held the multiplier neutral.

LIFTED BY moderate interaction-design signals; operators favour brief, task-focused sessions, which suits a purpose-built dashboard.

HELD BACK BY no satisfaction or feedback evidence at all; engagement signals very low.

TO RAISE CONFIDENCE structured user research with pilot-site operators; baseline usability metrics.

SECTION 05 · THREE GATES BEFORE THE €150M

Conditions, not milestones

Every dimension card on the previous page ends the same way: confidence rises only with evidence a pilot creates. The gates are that evidence, in order. These gates are not a project plan. Each one is a falsifiable condition with a defined consequence, and the first is a kill test: it carries a structured stop, with the decision date on the calendar before the first euro of scale capital moves. Notice their order. The contract is tested first, the price second, and the hardware only last: the analysis sequences product proof last because in this market it is the offer's shape, not its specification, that fails first. Pass all three and the €150M commitment proceeds on evidence. Fail the first and the capital stays in the core business, which the brief itself names as the alternative.

GATE 01

90 DAYS | REVERSIBLE | KILL TEST

Value proposition validation

PROBLEM	The TCO and ROI of the system are theoretical. Every commercial claim DataCore makes is currently a projection.
THE GATE	**Secure a signed, paid pilot agreement with a target hyperscaler.**
DO NOT YET	Commit to full-scale commercialisation or marketing spend.
UNLOCKS	Empirical proof of the value proposition and a reference customer, the single asset this market's conservatism demands.
MEASURE	Signed, paid pilot contracts. One is a pass. Zero at the deadline is the answer.
KILL IF	*"If no paid pilot can be signed, the value proposition is invalid." The sprint is set at 90 days, and the overall verdict flips to no-go if no hyperscaler will agree to a paid pilot within two quarters.*

PAID MATTERS

A free pilot tests the technology; only a paid one tests the proposition. In a market whose strongest competitor is a contract, the signature is the data: it is the one piece of evidence the run's thirteen unresolved tensions cannot supply and cannot survive.

ON FEASIBILITY

Name the hardest tension in this report rather than carry it silently: the same evidence that says this buyer procures slowly says a signed, paid hyperscaler pilot in 90 days is improbable by the direct route. The gate is built for that. The kill threshold is two quarters, not ninety days; doors B and C on the next page exist because they are faster; and the direct door opens mainly where grid distress has already overridden procurement pace. The sprint is deliberately set faster than the market's habit, because the habit is what is being tested.

SECTION 05 · THE GATES — CONTINUED, AND WHERE TO AIM THE PILOT

GATE 02 — 6 MONTHS

Business model viability

PROBLEM	The storage-as-a-service model's financial viability is unproven, and Section 02 shows a buyer that resists purchasing any other way.
THE GATE	**The financial model and TCO calculations validated by a target customer's finance and procurement team.**
DO NOT YET	Build out a sales team on the service model.
UNLOCKS	A bankable commercial model that accelerates every subsequent sale.
MEASURE	Written sign-off on the TCO model from the customer's finance and procurement team, not a courtesy review.
INVERSION	*"If the TCO is not compelling to finance buyers, the business model fails."*

GATE 03 — 9–12 MONTHS

Whole product completion

PROBLEM	The current product is incomplete for a pragmatist buyer; capability statements are the client's own, with no deployment behind them.
THE GATE	**The pilot completed, meeting all technical and operational criteria defined by the customer.**
DO NOT YET	Expand the product roadmap with new features.
UNLOCKS	A market-ready whole product and the case study that converts the early majority.
MEASURE	The buyer-journey analysis sets the bar itself: pilot system uptime meeting the Tier III target of 99.982%, alongside the customer's own acceptance criteria.
INVERSION	*"If the product fails to meet pilot KPIs, it is not ready for market."*

THREE PLACES GATE 01 CAN BE WON. CHOOSE BY INTERNAL FIT.

A · A grid-constrained hyperscaler, direct

Target an early-adopter hyperscaler building in a region with long interconnection queues, where speed-to-power is worth paying to test. Strongest reference value; slowest door to open.

B · Through a partner's trust

Approach an established industrial energy provider and run the pilot for one of its existing customers, borrowing credibility Voltaxis has not yet earned in this channel. Fastest path past the brand gap; shares the reference. The June 2026 verification pass adds urgency here: specialist storage developers are already signing the first data-centre deals, which makes them partners to approach now or rivals to meet later.

C · A large colocation operator first

A colocation buyer carries lower procurement gravity than a top-five hyperscaler and diversifies the customer base from day one. Easier signature; lighter halo.

TARGETING NOTE One trend widens the map for all three directions. The IEA expects accelerated-server capital to rise from 10% of data-centre investment in 2015 to 45% by 2030 as greenfield building slows, which shifts demand toward retrofitting existing sites, exactly where a modular behind-the-meter system beats a new grid connection on time.

The analysis defines the twelve-month success milestone in a single sentence: a paid pilot converted into a public, referenceable case study with a globally recognised hyperscaler. It also flags the machinery that does not yet exist anywhere in the plan: no programme converts a successful pilot customer into a public advocate. Build that programme during the pilot, not after it. In a market where the first contract is the position, the reference is the product of the year. One sequencing note: the gates nest rather than queue. The financial model that clears Gate 2 is built inside Gate 1's negotiation, and Gate 3 is the pilot those two gates created, run to completion.

SECTION 06 · THE BUYING TABLE

Three buyers, one signature each

A hyperscaler pilot is signed three times. The buyer analysis resolves the committee into three archetypes; the pilot offer has to clear each one in their own language.

BUYER 01 — CFO / VP OF INFRASTRUCTURE

The Economic Buyer

PROFILE	Approves large-scale expenditure on financial return and risk mitigation. Primary drive: reduce total cost of ownership and shift infrastructure cost from capital to operating expense.
JOB TO BE DONE	*"Help me be seen as making innovative but financially sound infrastructure decisions."*
SAYS YES WHEN	A validated financial model shows clear multi-year TCO savings against utility power and diesel backup.
LEAD WITH	Financial predictability and risk reduction.
AVOID	A capital-purchase framing; this buyer's aversion to owning power assets is the market's defining friction.

BUYER 02 — DATA CENTRE OPERATIONS DIRECTOR

The Technical Buyer

PROFILE	Evaluates reliability, integration, and operational impact. Primary drive: guarantee uptime and remove grid-instability risk.
JOB TO BE DONE	*"Give me confidence that I am not risking my career on unproven technology."*
SAYS YES WHEN	A pilot demonstrates utility-grade reliability and seamless integration without disrupting live operations.
LEAD WITH	Operational resilience and utility-grade safety.
AVOID	Leading with AI autonomy. The evidence shows operators distrust AI in system-critical roles; position the software as decision support under human control, and stage its autonomy deliberately. The journey evidence shows trust here is earned in steps, from assistant to automated arbitrage, never granted at installation.

BUYER 03 — HEAD OF SUSTAINABILITY

The Strategic Buyer

PROFILE	Ensures infrastructure aligns with public ESG commitments. Primary drive: meet 100% renewable targets and improve the facility's environmental profile.
JOB TO BE DONE	*"Help me be seen as ahead of the energy transition, not chasing it."*
SAYS YES WHEN	Third-party certification shows the system measurably reduces carbon footprint and fossil reliance.
LEAD WITH	A credible, certified path to stated sustainability mandates.
AVOID	Uncertified green claims; the analysis flags greenwashing exposure as a brand fragility for a new entrant.

SECTION 07 · MARKET SIGNALS AND THE COMPETITIVE FRAME

The strongest competitor is a contract. The second is waiting.

DEMAND BY 2030

·945 TWh

Global data-centre electricity by 2030 in the IEA base case, more than double 2024's roughly 415 TWh.

ADDRESSABLE PAIN

·20%

Share of global data-centre capacity planned to 2030 that grid constraints could delay, per the IEA. This is DataCore's addressable pain.

CAPITAL IN MOTION

$300B

AI-related capital expenditure the four largest hyperscalers were reported to be planning for 2025 alone.

Buying power is concentrating as it grows: hyperscale facilities have moved from roughly 10% of data-centre capacity in 2010 to 37% today. A small set of very large accounts now drives the demand curve, which cuts both ways. One signed hyperscaler validates the category; none signed is a verdict. Behaviour signals from adjacent industrial buyers say the software layer lands on prepared ground. In a 2018 survey of 157 European manufacturers, 68% used energy-management dashboards daily, though that is a proxy from a neighbouring market, not data-centre evidence.

THE REAL COMPETITION

Ranked by value to the buyer, the analysis puts hardware last.

ALTERNATIVE	WHY BUYERS CHOOSE IT	ITS WEAKNESS
Corporate power purchase agreements	No assets on the buyer's balance sheet; ESG-friendly	Inflexible long terms; does not solve on-site resilience or interconnection delay.
Do nothing, delay the build	No capital outlay, no new risk	Pays the opportunity cost of stalled AI capacity while the queue does not move.
Grid expansion	The default; zero operational change	Interconnection queues measured in years; the exact constraint driving this market.
Conventional UPS (lithium-ion)	Established, trusted, certified	Minutes of cover, not the 8–12 hours of sustained resilience the client specifies for AI loads.
On-site gas generation *(added at verification, June 2026)*	The fastest primary-power fix now being deployed at scale	Cannot provide the instantaneous response, load smoothing and clean firming that volatile AI loads and interconnection deals demand; batteries are being positioned as its complement, not its casualty.

SECTION 07 · THE COMPETITIVE FRAME — CONTINUED

The clearest emblem in the evidence base is a dispute. Microsoft publicly challenged Georgia Power's load-forecasting methods, arguing the utility undervalues renewable energy and overestimates data-centre load for projects that have not chosen it as a provider. When the analysis names its strategic enemy as the status quo of utility dependence, this is what that looks like in the wild. The buyer-utility relationship is fraying in public, and every fray is an opening for a seller offering control on the buyer's side of the meter.

Named manufacturers, per the brief, include Tesla, Fluence, Wärtsilä, Sungrow, CATL, LG Energy Solution, Powin, and hyperscalers' own in-house teams; industry analysis names Tesla, Fluence and LG Energy Solution as key players in US data-centre battery storage. None yet owns the AI-native, behind-the-meter, service-structured position DataCore is aimed at, and Brand Finance reports the most visible name on that list, Tesla, weakening markedly in Europe.

VERIFIED AGAINST THE MARKET, JUNE 2026

Attributed external context; the evidence base of the analysis predates it.

The evidence base behind this analysis closes in early 2025. A verification pass against the market as of June 2026 strengthens the thesis and shortens the clock. Grid queues in primary data-centre markets now exceed four years, and behind-the-meter arrangements have moved from concept to the industry's stated playbook. The first data-centre battery deals are public, including a colocation operator commissioning a 31 MW on-site battery to connect years earlier than the grid would allow. Industry surveys still rank power purchase agreements as the most common strategy, with on-site generation third and rising, and operators describe running power assets as a business they are not built for. A minority counter-current is real and worth watching: a few of the largest AI builders are moving toward owning generation outright, though the deal-structure evidence currently outweighs it. **Both halves of the spine hold: the buyer still solves power through contracts, and the contracts are now being signed.** What has changed is the tempo: the most visible battery name on the named-competitor list is already selling into this space, and storage developers are signing the early deals. The 90-day clock in Section 05 is not a planning convenience. It is the market's own speed.

FIVE FORCES

Levels only — directional model output.

BUYER POWER	SUBSTITUTE PRESSURE	RIVALRY	NEW ENTRANTS	SUPPLIER POWER
High	Medium	Medium	Low	Low

Supplier power carries one named exception: China refines around 99% of global gallium, and data-centre demand for it could exceed 10% of today's supply by 2030, a watch-item for power-electronics sourcing.

The walls that keep new rivals out do not make the buyer sign. The pressure that does is arriving through the gap between the buyer's preferred contract and the demand it cannot cover.

SECTION 08 · HOW THE NEXT TWO YEARS COULD BREAK

Two futures are tied. The plan must survive the harder one.

RISK-AVERSE MARKET · 39%

Risk aversion dominates buying. New technology is judged on perceived likelihood of failure, and incumbents are the safe choice. The service model and a reference customer stop being advantages and become entry requirements.

HOLDS IF · buyers prioritise cost discipline over innovation, and known vendors keep the benefit of the doubt.

AGGRESSIVE GROWTH · 39%

Demand outruns supply and speed becomes the only differentiator that matters; the constraint shifts to supply chain and deployment capacity, and the fastest deployer wins the largest contracts.

HOLDS IF · AI build-outs continue to outpace grid response, and deployment speed outweighs marginal cost differences.

DIFFERENTIATION UPSIDE · 22%

Specialisation is rewarded. Purpose-built systems for AI loads command premium pricing among buyers who feel the volatility pain most acutely, and generalist storage reads as inadequate.

HOLDS IF · buyers can distinguish specialised performance and will pay for it.

TWO STRUCTURAL RISKS — EACH WITH ITS BUILT-IN DEFENCE

Risk 01 · Capital burn without traction. The first named risk: resources flow into fixing one of the weak dimensions, product completeness, brand presence, experience evidence, before validation confirms it is the bottleneck a paying customer actually cares about. **The defence is the sequence itself:** the do-not-yet freezes in Section 09 and the gate order ensure no scale spend precedes a signature.

Risk 02 · The demand signal is a mirage. The second named risk: the moderate market reading could mask the absence of urgent, addressable pain in any single buyer profile, so even a perfected product finds no signature, a risk amplified by the run's moderate-to-low confidence. **The defence is the kill test:** a paid pilot is the only instrument that resolves this question, and the 90-day gate surfaces the answer early instead of absorbing it late.

The two risks reinforce each other, and the gates are how the two favourable scenarios become available rather than hypothetical.

READ THIS BAR BOTH WAYS

The bar's first segment carries the model's dominant tag even though growth ties it at 39%. Read it as a statement about which risk to plan for, not which future is likelier. The adverse case is itself tagged balanced risk-reward rather than catastrophic. The strategic instruction is identical under both: in a risk-averse market the paid pilot is the entry ticket, and in a growth market it is the head start. The gates are robust to the tie. Probabilities are directional model output, not forecasts.

SECTION 09 · THE NEXT TWELVE MONTHS, IN ORDER

Spend nothing that scale requires until validation pays for it

DO FIRST

Months 0–3 — the resource plan the analysis itself specifies.

Run targeted outreach to five hyperscaler chief technology offices, a sniper approach aimed at one yes, not a broadcast; if direction B or C on page 9 fits the company better, the same discipline applies to that door. Lead with guaranteed power resilience for mission-critical AI delivered as a service; support with a TCO model priced against the buyer's true alternative, grid power plus interconnection delay plus conventional backup; close on a limited-slot, paid pilot. In parallel, finalise the service financial model to the standard a customer's finance team can interrogate; Gate 2 begins inside Gate 1's first meeting. Dedicate the company's best engineering talent to pilot support, because a single technical failure can destroy the reference asset everything else depends on.

DO NOT YET

Verbatim substance from the analysis.

No full-scale commercialisation or marketing spend. No sales-team build-out on the service model. No product-roadmap expansion beyond what the pilot's acceptance criteria require. The posture is **hold expansion**: every euro ahead of evidence is the capital-burn risk the analysis flags as this venture's primary failure mode.

THE SEQUENCE — TWELVE ACTIONS

Months as order-of-magnitude. All timings estimated. Each chip carries the gate it serves.

Now
MONTHS 0–1

- Deploy initial hyperscaler pilot system — GATE 1
- Build the customer-grade TCO and ROI model — GATE 2

Short
MONTHS 2–4

- Tier III-compliant redundancy engineering — GATE 3
- Predictive load-management software hardening — GATE 3
- Operator dashboard, human-in-control framing — GATE 3
- Speed-to-power go-to-market materials for pilot targets only — GATE 1

Mid
MONTHS 5–7

- Analyst and trade-press credibility programme — SCALE PREP
- Real-time TCO and arbitrage reporting in product — GATE 2
- Remote diagnostics and NOC integration — GATE 3

Long
MONTHS 8–11

- Data-centre design-firm channel partnership — SCALE PREP
- Storage-as-a-service commercial launch, *gated on Gate 2* — SCALE
- Supply-chain diversification plan, including gallium-exposed power electronics — OPEN QUESTION 3

VALUE NOTE The analysis also prices these actions, and the pricing confirms the order: the pilot itself is the single highest-value move on the board, with the TCO model, the reliability engineering and the eventual service-model launch ranked just behind it. The sequence is a value ranking as much as a calendar, and nothing rated highly sits outside a gate's service.

ECONOMICS NOTE The baseline economics in this analysis are synthetic, built on proxy unit assumptions rather than Voltaxis cost data, and they read as cost-driven: attractive only under volume assumptions that validation has not yet earned. Treat unit economics as an open question for Gate 2, not an input to the decision.

SECTION 10 · OPEN QUESTIONS

The questions that decide more than the scores do

Q1 What is the quantifiable TCO advantage of the service model for a typical hyperscaler data centre?

WHAT WOULD CHANGE A proven, positive TCO unlocks the sales process; a negative or marginal one invalidates the core business model.

EVIDENCE NEEDED A completed financial model validated by a target customer's procurement team during pilot negotiation.

Q2 Which specific operational risks do data-centre operators fear most that on-site storage can solve?

WHAT WOULD CHANGE Aligning the proposition with the top one or two perceived risks would materially shorten the sales cycle.

EVIDENCE NEEDED Direct interviews with five to ten data-centre operations directors.

Q3 Can Voltaxis build a supply chain that scales to hyperscaler deployment timelines?

WHAT WOULD CHANGE Confirmation makes Voltaxis a credible long-term partner; failure confines it to small-scale niche work whatever the pilot proves.

EVIDENCE NEEDED A feasibility study on vertical integration or strategic partnership for key components.

Q4 Is 8 to 12 hours the right size for the firming slot, or is the defensible product a shorter system with the software doing the differentiating?

Raised at the June 2026 verification, not by the baseline analysis.

WHAT WOULD CHANGE The first deals in this slot run on two-hour systems doing the smoothing and interconnection jobs this report assigns to DataCore. If the slot wants short duration, the client's central specification is oversized for it, with the cost structure that implies; if customers will pay for sustained ride-through, the longer duration becomes the differentiator the brief assumes.

EVIDENCE NEEDED The pilot customer's own acceptance criteria and the duration they are willing to price, discovered inside Gates 1 and 2.

SECTION 10 · OPEN QUESTIONS — HOW THEY RESOLVE

What could still move the verdict, and what cannot

The four questions do not all resolve on the same clock, and none of them unsettles the verdict. This page sequences them, names the one condition that would flip the decision, and names where the capital goes if it does.

HOW THESE CONNECT

The second question is first-90-days work: ranked operator fears shape the pilot offer itself. The first and the fourth resolve inside Gates 1 and 2, not before them; willingness to pay, and the duration it attaches to, are discovered in negotiation, not in research. The third is the 6-to-12-month strategic resolution that decides how large a cleared position can become. One adjacency the run could not see from public data: the 2026 deals are being struck through storage-developer and utility partnerships, which raises an internal question worth answering in the first month: which of Voltaxis's existing utility relationships already open data-centre doors. The verdict holds regardless of how these resolve. What shifts is the execution shape, and the gates are built to absorb that.

INVERSION The verdict flips to no-go if no hyperscaler will agree to a paid pilot within two quarters, signalling the value proposition is not compelling.

IF THE GATE FAILS A no-go is not a dead end, and the closing line's "or pivot" points at something specific. The capital stays in the utility-scale core, which the brief itself names as the alternative, and the analysis names one adjacent redeployment of the same technology worth a desk review in that event: pointing the load-intelligence and scalability stack at grid operators themselves, a market the analysis reads as less contested and carrying a more urgent version of the same problem. A failed Gate 1 would invalidate the data-centre proposition, not the stack.

This assessment is based on public and proxy inputs. Additional internal data would materially improve precision and relevance.

SECTION 11 · METHOD & SOURCE BASIS

A structured decision intelligence analysis of the Voltaxis DataCore market-entry question, evaluated across five dimensions (Market, Consumer, Product, Brand, Experience). Twenty-eight analytical factors were scored from 143 evidence items across 111 sources, weighted by source quality and cross-stream agreement; the thirteen places where the evidence streams could not be reconciled were penalised in the scores rather than smoothed over.

UPLOADED SOURCE DOCUMENTS *client-supplied primary evidence*

LBNL, US Data Center Energy Usage Report · Dec 2024	US consumption baseline and projections (176 TWh, 4.4%, 325–580 TWh by 2028); growth-rate history; the Microsoft–Georgia Power dispute; reference chain to Goldman Sachs.
IEA, Energy and AI · WEO special report, 2025	Global demand outlook (945 TWh by 2030); accelerated-server growth rates; the asset-ownership keystone sentence; PPA pipeline coverage; hyperscaler capex; gallium exposure; grid-delay risk.

KEY CITED SOURCES *engine evidence base*

Goldman Sachs Research · May 2024	160% data-centre power demand growth projection to 2030.
Energy Policy; Energy and Buildings · 2016 · 2018 · 2019	Industrial energy-management behaviour proxies — dashboard use, provider renewal, interface engagement — each anchored to its survey sample and year in the text.
Uptime Institute	Tier III availability target (99.982%).
MarketsandMarkets · 2024	Named key players in US data-centre battery storage.
Brand Finance · 2025–2026	Tesla brand-strength trajectory in Europe, publicly reported.

JUNE 2026 VERIFICATION SOURCES *external context, separate from engine evidence*

JLL, Global Data Center Outlook · 2026	Grid queues beyond four years; behind-the-meter expectations.
Latitude Media · Nov 2025	First data-centre battery deal of its kind (31 MW, developer-built, earlier interconnection).
Foley Data Center Survey, via datacenterHawk · 2026	Power-strategy rankings (PPAs first; on-site third).
Bank of America Securities, via Utility Dive · Feb 2026	The 2026 power hierarchy: gas first, storage as firming.
pv magazine USA; 451 Research; Energy Storage Summit USA · Mar–Apr 2026	Batteries as the firming complement; operator skills gap; behind-the-meter consensus.

A NOTE ON THE SCORES

The dimension scores, the 53/37 readiness and differentiation readings, and the 39/39/22 scenario distribution are **directional analytical outputs, not forecasts**. Confidence values run 0.0 to 1.0 and measure data depth, not certainty; this run's evidence confidence is moderate to low throughout, disclosed in Section 01 and reflected in the conditional shape of the verdict.

The engine's synthetic unit economics were excluded from this report; commercial figures await Gate 2. No product documentation was part of the run, so all product capability statements are the client's own.

CLOSING NOTE

The brief posed a binary: invest €150M or stay focused. The evidence supports neither answer as posed. It supports a sequence: a capped 90-day sprint to a signed, paid pilot, two further gates that convert theory into bankable proof, and a standing instruction that the scale capital moves only when the gates do. Underneath the sequence sits the finding the whole analysis turns on: **in this market the first signed, paid contract is not a sale. It is the position, and the moat that forms around it.**

PRELIMINARY READ

"The immediate imperative is to convert market potential into market proof; secure a paid pilot within 90 days to validate the value proposition or pivot."

VALIDATUS | VOLTAXIS

Decisions you can defend.