White Paper

The Commercial & Industrial Distributed AI Data Center

Turning behind-the-meter battery sites into a distributed AI compute cloud — at industrial scale.

Built on JVPower.AI industrial construction

Henry Pham | President & CEO
JVPower.AI | Katy, TX 77494
hpham@jvpower.ai | +1 (281) 948-6878 | jvpower.ai

Robert Sliva | Vice President & COO
JVPower.AI | Katy, TX 77494
bsliva@jvpower.ai | +1 (713) 857-7321 | jvpower.ai

Contents

01

Executive Summary

The world's demand for compute is surging, while the power to serve it is increasingly constrained.

JVPOWER already does the hard, capital-intensive work: it finances, builds, and operates behind-the-meter battery energy storage (BESS) on commercial and industrial (C&I) properties, with deep construction execution inherited from JVBUILD. This white paper sets out how the same sites, the same electrical infrastructure, and the same construction capability can carry a second, far larger revenue stream — distributed AI compute — turning every C&I property into a node that earns from both energy and inference.

JVPOWER is the frontier of commercial- and industrial-scale distributed AI compute. The company deploys dispatchable 20- and 40-foot containers holding hundreds of GPUs onto C&I sites where the power, the interconnection, and — critically — the battery already exist. Each pod taps hundreds of kilowatts of spare capacity behind a single industrial service that already carries JVPOWER storage, then coordinates many sites into one low-latency compute cloud. This is purpose-built for industrial scale: large services, high load factors, and fewer, bigger deployments mean far lower cost per megawatt and a clean, commercial host relationship.

176→580	2.7×	~49 GW	4.5–7 yr
U.S. data-center TWh, 2023→2028	global DC demand, 2025→2030	projected U.S. power shortfall	grid interconnection wait

How JVPOWER makes money, in one line: own the container + cooling + power + BESS, then either lease compute-ready capacity to a creditworthy AI customer (a ~3-year payback, ~30%+ annual cash yield, low-risk infrastructure asset) or — against a contracted offtake — also own the GPUs and sell inference at ~80% gross margin. How JVPOWER pays the host: guaranteed lower energy cost, a monthly hosting payment, backup resilience, and zero capital outlay. Both sides are quantified in Sections 8–10.

02

Macro Challenge: Compute vs. Power

2.1 Exploding compute demand

AI has broken a decades-long flat trend in electricity demand. U.S. data centers consumed about 176 TWh in 2023 (~4.4% of national consumption) and are projected to reach 325–580 TWh by 2028 — a doubling-to-tripling in five years. Globally, McKinsey models data-center power demand rising from ~82 GW in 2025 to ~219 GW by 2030, with AI inference — latency-sensitive and best served close to users — overtaking training to become the single largest workload.

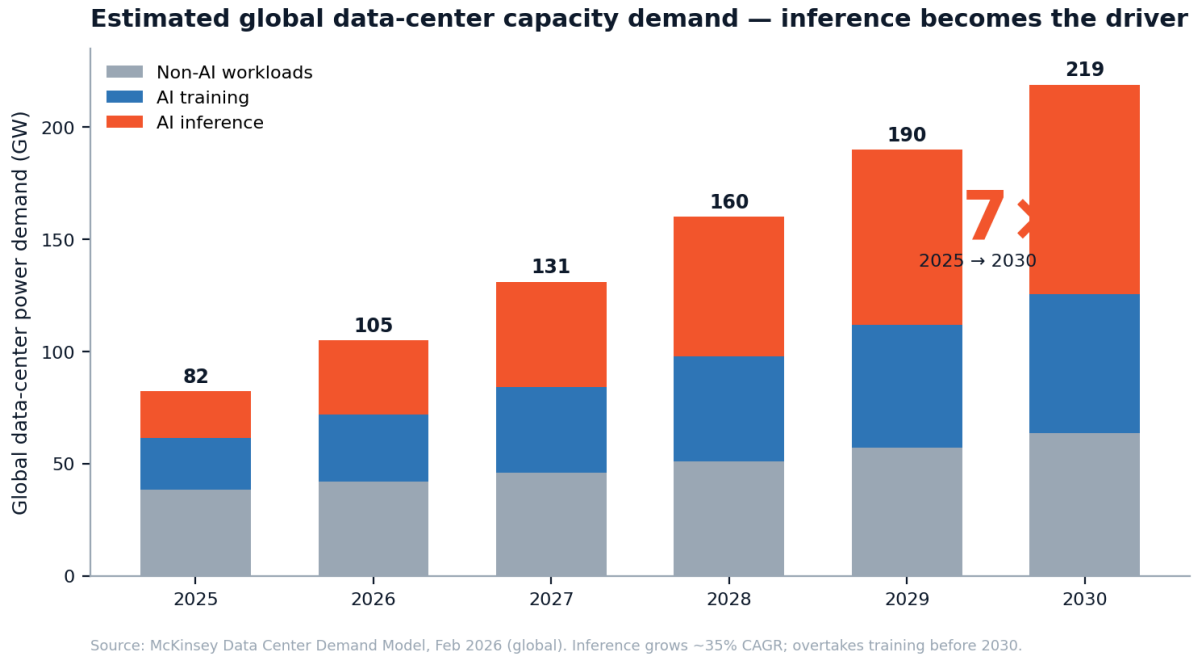

Global data-center capacity demand grows ~2.7× by 2030, led by inference. Source: McKinsey Data Center Demand Model, Feb 2026.

2.2 Grid capacity & local interconnect limits

The binding constraint is not chips or capital — it is power delivery. Roughly 2,290 GW of capacity sat in U.S. interconnection queues at the end of 2024, nearly twice the installed fleet, and projects reaching operation that year had waited about 4.5 years on average; in the hottest markets, large-load interconnection stretches to 7–10 years. A widely-cited analysis puts the near-term U.S. generation shortfall at roughly 49 GW by 2028. This is the “speed-to-power gap”: demand arriving years ahead of the grid’s ability to serve it.

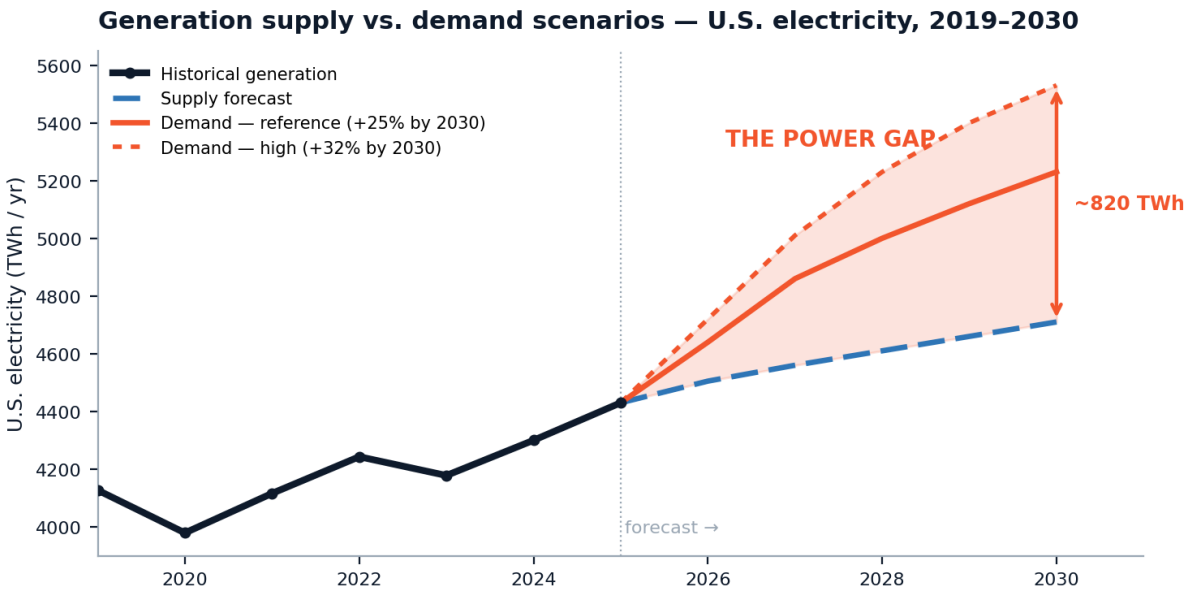

Source: EIA (historical & supply forecast); Grid Strategies & Pew/McKinsey (demand scenarios). Planning estimate, 2026.

Supply forecast vs. demand scenarios for U.S. electricity. The widening gap is the opportunity JVPOWER addresses with already-energized sites.

Two shifts make distributed, behind-the-meter compute the natural response. First, inference rewards proximity to users. Second, large loads can be made flexible — a Duke University study found the existing grid could absorb ~100 GW of new load if it accepts brief curtailment, and regulators are codifying this (NERC's 2025 large-load recommendation; Texas SB6). A battery-buffered, curtailable container on an already-powered C&I site is exactly the asset this moment calls for.

03

What is JVPOWER?

A distributed compute cloud built from a fleet of AI containers on commercial & industrial sites.

JVPOWER converts underutilized behind-the-meter power capacity at C&I properties into high-value AI compute. Each site pairs a factory-built, liquid-cooled GPU container with JVPOWER's on-site battery and power-control electronics; the containers are coordinated through an orchestration layer (JSOL) so that many independent nodes behave like one coherent, low-latency inference cloud. JVPOWER does not replace hyperscale data centers — it augments them by adding dispatchable capacity at the grid edge, near where inference demand actually lives.

A JVPOWER node on a commercial/industrial site: the AI container plus on-site BESS and dry coolers, beside the host facility.

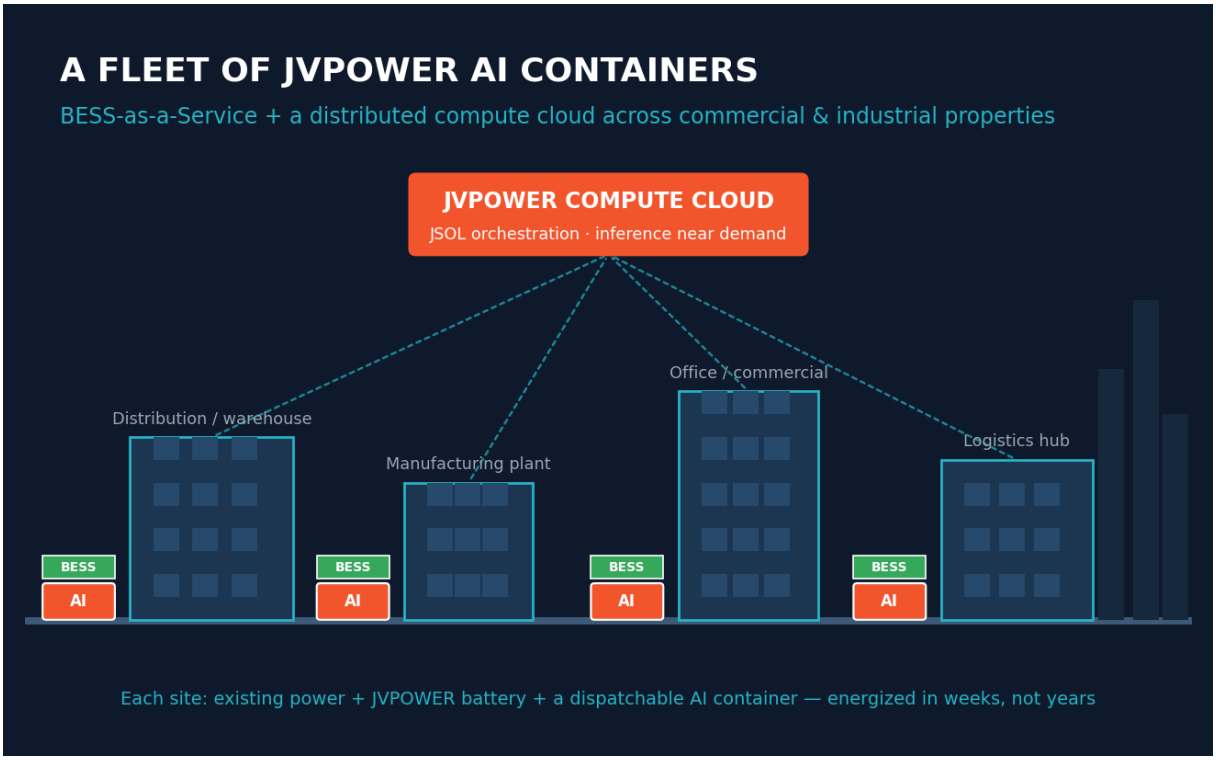

A fleet of JVPOWER AI containers across C&I properties: BESS-as-a-Service plus a distributed compute cloud, energized in weeks rather than years.

04

The C&I Unlock: Battery + Power Control

4.1 Unlocking existing industrial headroom

Industrial electrical services are sized for rare peak conditions and typically run far below their rated capacity. Under NEC 220.87, existing service can be re-rated against measured peak demand rather than conservative nameplate calculations — and a Stanford grid-utilization study found that even at annual peak, utility transformers ran at only ~28–68% of capacity (most ~45%). That spare capacity, behind a single large industrial meter, is JVPOWER's raw material.

4.2 Making headroom usable for always-on compute

Headroom only becomes useful if compute can run safely and predictably. JVPOWER's battery is the unlock. GPU loads swing from near-idle to full draw in milliseconds; the BESS absorbs those spikes on-site so they never hit the meter or the grid, shaves the host's demand peaks, and — paired with a UL 3141-class power-control system — lets the container be added without a utility service upgrade or a place in the interconnection queue. The same battery makes the load curtailable in seconds, satisfying emerging flexible-load rules and creating grid-service value.

4.3 Rapid deployment at scale

Because JVPOWER deploys behind the existing meter, and because the containers are factory-built and pre-integrated, nodes drop in incrementally — modular data centers come online in roughly 6–9 months versus years for centralized builds. JVBUILD's construction capability turns the hardest part of this market (siting, MEP, interconnection, install) into core competency. Capacity is added across a fleet and coordinated through JSOL as a single distributed cloud.

05

Addressing the Big Questions

5.1 Uptime, reliability & continuity

Containers are engineered for enterprise-grade availability. The on-site BESS provides ride-through and orderly shutdown during grid events; during a utility demand-response or outage event, JSOL fails workloads over to unaffected nodes elsewhere in the fleet, and the battery powers a graceful pause or migration. Distributed capacity means no single site is a single point of failure.

5.2 Sustainability, water & grid impact

JVPOWER raises utilization of existing grid assets rather than forcing new generation and transmission, and the battery acts as a grid “shock absorber” during stress events. Cooling is fully integrated inside the AI container — closed-loop direct-to-chip liquid through a CDU to a dry/fluid coil and roof-mounted EC fans, engineered for a PUE of 1.05. There is no external cooler, chiller, or cooling tower; ongoing water use is essentially zero, with adiabatic assist drawing on a backup water line only on the hottest days. That makes the container quiet, dry, and easy to site — a decisive advantage in drought-prone regions and for permitting.

5.3 Host value & fairness

The C&I property owner gets a compelling deal (detailed in Section 8): a contractually lower energy cost, a monthly hosting payment, backup resilience for their own critical loads, and zero capital outlay — JVPOWER finances, builds, owns, and operates everything. When a node must throttle for a grid event, the host’s power is protected first.

5.4 Security & trust

Containers are secure, monitored enclosures with encrypted communications, tamper-evident protections, and full workload isolation suitable for multi-tenant compute. Physical and cyber security are designed for unattended operation across a distributed fleet; compromised hardware can be remotely revoked and disabled.

06

Target Workloads

Best suited to workloads that run on smaller hardware instances, benefit from incremental capacity, and value low latency.

JVPOWER is a strong complement to large data centers, serving demand closer to end users while expanding capacity faster than centralized build-outs alone. It targets today's highest-growth workloads:

35%	34%	8%	7%
AI inference CAGR	Cloud gaming	Cloud rendering & streaming	HPC & simulation

For interactive and inference workloads, proximity to users directly improves experience and reduces backbone congestion — precisely where distributed C&I nodes outperform remote mega-campuses.

07

JVPOWER Concept & Architecture

7.1 Container node architecture

A site is two containers: a BESS container and an AI Data Center container. Inside the AI container, the chain is power (from the BESS bus) → switchgear and power-control system → liquid-cooled GPU racks (~56 kW/rack) → coolant distribution unit → an integrated dry/fluid coil and roof-mounted EC fans for heat rejection, coordinated by JSOL at PUE 1.05. The battery firms the GPU load and provides backup; cooling is entirely inside the container, so the only site connections are power and a backup water line.

Two-container site: a BESS container and an AI Data Center container with cooling fully integrated inside (CDU + dry/fluid coil + roof EC fans). No external cooler — only two site connections: power and a backup water line.

Three reference sizes for individual pods: a 20-ft pod (384 GPUs, ~334 kW IT, 800 A service), a 30-ft pod (576 GPUs, ~501 kW IT, 1,000 A), and a 40-ft pod (768 GPUs, ~668 kW IT, 1,600 A) — all at PUE 1.05. A 20-ft pod fits within spare capacity many C&I sites already have; larger pods use the BESS and power control to add load without a utility upgrade. Large C&I campuses — industrial parks, logistics REITs, and heavy manufacturing facilities — support a ten-pod cluster deployment: ~7,680 GPUs, ~6.68 MW IT load, ~16,000 A service, delivering 10× the compute density and revenue of a single-pod site at the same unit economics and the same two-site-connection footprint. Each pod connects to the site through three pre-engineered interfaces — a power tie-in from the

BESS bus, a redundant fiber uplink, and a backup-water line — bundled as a Site Connection & Service package that adds only about 1–2% to site capital cost (roughly $0.12M–$0.24M per pod; itemized in the JVPOWER electrical drawing set, DWG-E6 Rev 6).

7.2 Fleet orchestration & power events

At the fleet level, JSOL schedules and routes workloads across nodes based on customer requirements, latency, and real-time energy and grid constraints. During a grid event, JSOL surfaces the constraint, fails workloads over to other nodes, and powers affected nodes down safely on battery — delivering managed, predictable behavior that puts host and grid safety first.

7.3 Why this architecture matters

Together these elements create an infrastructure layer that delivers compute and power at the edge, reduces reliance on slow centralized build-outs, and increases resilience by distributing capacity across many independent, already-energized deployments.

08

Core Value Propositions

JVPOWER turns underutilized C&I power into a managed, distributed compute layer — and shares the value with the property owner.

Revenue per megawatt is why stacking compute on an energy asset is so powerful: AI compute monetizes at ~$3M+ per MW-year, versus ~$1.2M for a conventional data center and only ~$0.04–0.11M for merchant battery grid services — roughly 30–80× the energy-only revenue per megawatt.

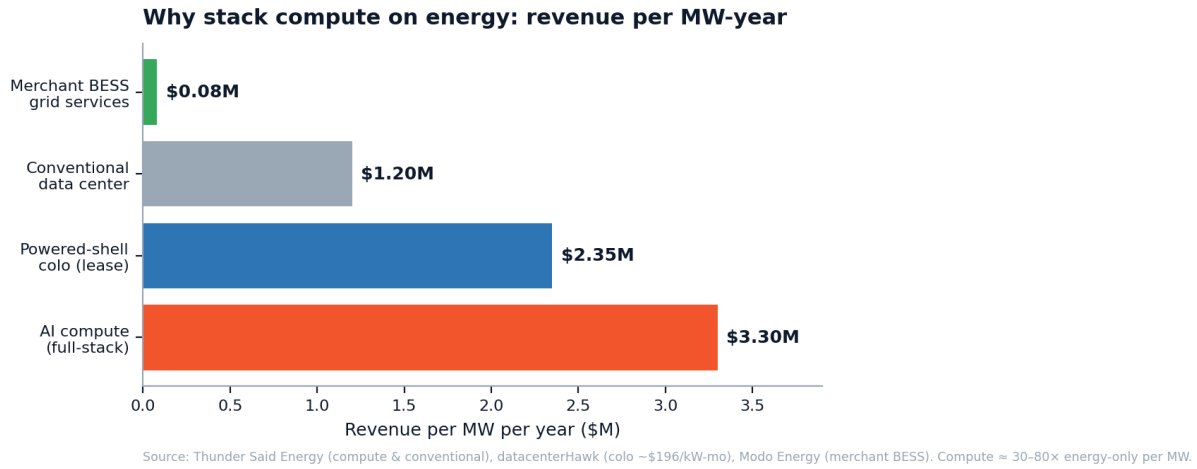

Revenue per MW-year by model. Sources: Thunder Said Energy; datacenterHawk; Modo Energy.

How JVPOWER makes money — two models:

Model A — Powered-shell colocation (recommended lead)

JVPOWER owns the container, cooling, power and BESS and leases compute-ready capacity to a creditworthy AI customer who brings the GPUs (~$180/kW-month). Low capex, low technology risk, recurring revenue: a single 20-ft pod nets ~$0.54M/yr on ~$1.7M capex — a ~3.2-year payback and ~32% annual cash yield on a 10-year asset. A large C&I campus running ten 40-ft pods scales to ~$5.4M/yr revenue on ~$17M capex at the same unit economics — the preferred deployment for industrial parks, logistics REITs, and heavy manufacturing sites with 10× the electrical headroom of a typical C&I property.

Model B — Full-stack compute

JVPOWER also owns the GPUs and sells inference directly: ~80% gross margin and ~2.5-year payback at year-one rates. Because merchant GPU prices fall ~25%/yr, this is run only against a contracted, take-or-pay offtake that locks rate and utilization — converting the headline margin into bankable cash flow.

How JVPOWER pays / charges the host:

The C&I host receives

(1) Guaranteed energy savings — a contractually lower effective energy cost of 15–30% via the BESS; (2) a monthly hosting payment — a lease or revenue share benchmarked at ~$10–$15 per kW of IT load per

month; (3) backup resilience for the host's own critical loads from the shared battery; and (4) zero capex and zero operational burden — JVPOWER finances, builds, owns, and runs everything. Cash lease plus larger energy savings, resilience, and no capital outlay make hosting a JVPOWER pod a straightforward decision for a commercial property owner.

Other stakeholders

Compute customers (hyperscalers, neoclouds, AI providers) gain flexible capacity without multi-year lead times. Utilities better manage peak demand and defer capital by leveraging existing, underutilized infrastructure. Investors fund durable, ITC-supported energy infrastructure with embedded compute upside.

09

Next Steps & Roadmap

Pilot (0–6 months). Deploy one 20-ft pod on an existing JVPOWER BESS site under a powered-shell agreement with a single anchor compute offtaker. Prove commissioning time, PUE, water use, uptime, and unit economics. Fund the storage layer with the 30% federal ITC (IRA §48E).

Anchor (6–18 months). Sign a take-or-pay offtake with a creditworthy customer (parent guaranty or letter of credit), standardize the host agreement (energy savings + hosting payment + curtailment clause), stand up a GPU financing SPV, and add the first 40-ft deployments.

Scale (18–36 months). Replicate across the JVPOWER / JVBUILD site portfolio and pipeline. Lock interconnection certainty under the new FERC co-located-load rules, finance pods on a repeatable infrastructure-plus-SPV stack, and grow JSOL into a single low-latency inference cloud. Partner with C&I property owners (e.g., Prologis-class logistics REITs) and industrial design-build contractors (e.g., Clayco, ARCO) to embed pods directly into new C&I construction.

10 Risks

JVPOWER approaches a new infrastructure model with disciplined risk reduction through staged pilots and measurable milestones. The chart maps the key risks by residual likelihood and impact after mitigation.

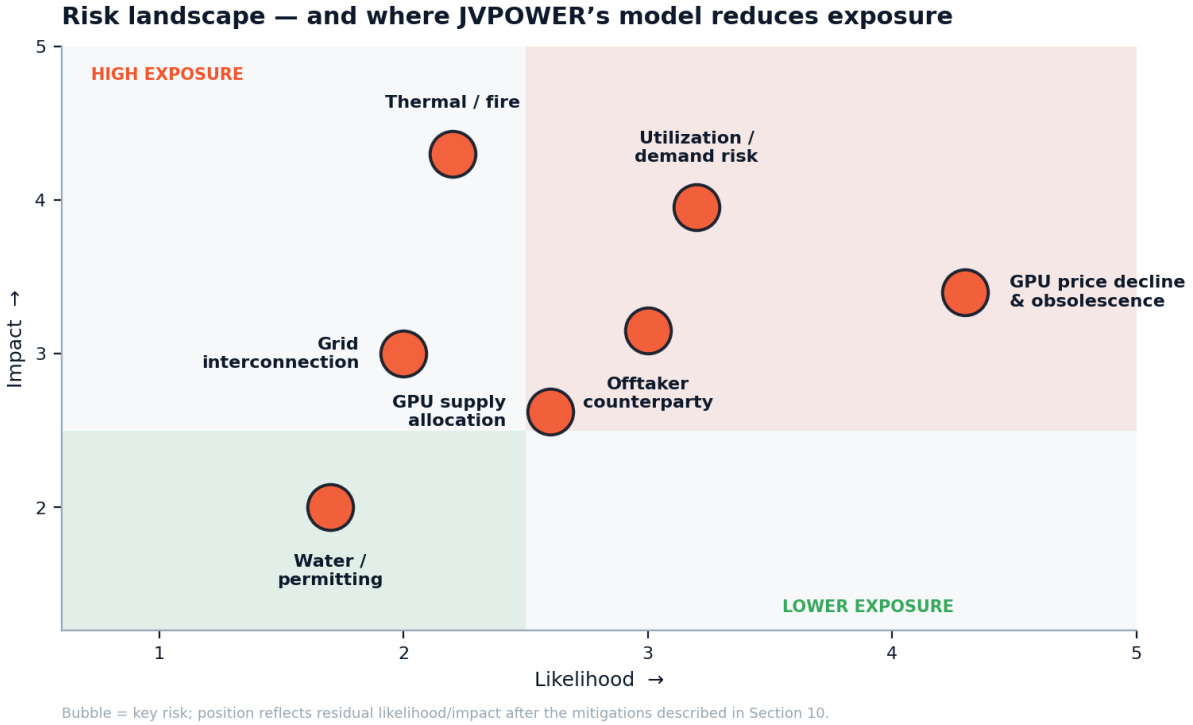

Risk landscape after the mitigations below. The model is deliberately structured to push exposure toward the lower-left.

GPU price decline & obsolescence — lead with powered-shell (the customer carries GPU risk); own GPUs only against take-or-pay; finance GPUs in a ring-fenced SPV sized to a 3–4-year life.

Utilization & counterparty — sign an anchor offtake before deployment; mitigate sub-investment-grade customers with parent guaranties, letters of credit, or hyperscaler wraps.

Interconnection, thermal & water — the BESS plus power control add load without an upgrade and make the load curtailable; closed-loop direct-to-chip cooling with NFPA 75 protection and dry coolers keeps water near zero.

11

Conclusion & Call to Action

JVPOWER is a pragmatic response to the growing imbalance between compute demand and power-delivery capacity.

By deploying dispatchable AI containers alongside the batteries JVPOWER already installs on C&I sites, the company converts underutilized behind-the-meter capacity into a managed, distributed compute layer — complementing centralized data centers while enabling faster, more geographically distributed expansion that aligns with grid realities. It earns from energy and compute at once, and pays the host property owner in cash, savings, and resilience.

Want to help build JVPOWER?

Investors — fund proof-of-concept execution and scale-out. Compute customers — anchor offtake for inference and cloud workloads. C&I property owners & builders — host nodes and share in the upside. Utilities — collaborate on grid-edge operations and demand response.

The demand is real. The grid constraints are real. JVPOWER is built to bridge the gap.

Contact

Henry Pham | President & CEO
JVPower.AI | Katy, TX 77494
hpham@jvpower.ai | +1 (281) 948-6878 | jvpower.ai

Robert Sliva | Vice President & COO
JVPower.AI | Katy, TX 77494
bsliva@jvpower.ai | +1 (713) 857-7321 | jvpower.ai

Sources

- Lawrence Berkeley National Laboratory / U.S. DOE, 2024 Report on U.S. Data Center Energy Use (Dec 2024).
- McKinsey, “The future of AI workloads,” Feb 2026 (global data-center demand model).
- EIA Short-Term & Annual Energy Outlooks; Grid Strategies, National Load Growth Report (Dec 2024).
- LBNL interconnection queue review (“Queued Up”), Dec 2025; Duke University large-load flexibility study, 2025.
- UL 3141 Power Control System standard (UL, 2024); 2026 National Electrical Code PCS provisions.
- Stanford / American Public Power Association transmission & transformer utilization study (Dec 2025).
- Thunder Said Energy (data-center & compute economics, 2025); datacenterHawk colocation pricing (2026); Modo Energy merchant BESS revenue (2025).
- IRA §48E storage ITC (IRS); FERC co-located-load order (Dec 2025); NERC large-load recommendation; Texas SB6 (2025).
- NVIDIA RTX PRO 6000 Blackwell Server Edition specifications; NREL ATB & BloombergNEF storage cost data (2025).

All figures are planning estimates as of June 2026. GPU pricing and AI-demand forecasts are volatile and should be treated as directional. Renderings and charts are conceptual/illustrative.