SATSIGHT · Grid-Interactive Infrastructure research

FLEXIBLE COMPUTE · VIRTUAL POWER PLANTS

Grids · Bitcoin mining · AI data centres

The Largest Battery the Grid Already Has

Data centres are no longer just the grid's largest loads — they are becoming its most valuable flexible asset.

A plain-language analysis of flexible compute as a grid resource — why Bitcoin mining has already proved the model, what AI workloads can and cannot inherit from it, and the Super-VPP architecture that combines them on one grid connection.

BITCOIN NETWORK	NORTH AMERICA	ERCOT · RIOT	US HEADROOM
~20 GW	**37.5 GW**	**$71M**	**98 GW**
continuous global load	VPP capacity behind the meter	grid credits paid in 2023	for load curtailing 0.5%/yr (Duke)

Prepared: 13 July 2026 · Companion to Satsight research note SSC-RN-2026-013 (published 8 June 2026) and the keynote film of the same title.

Sources: EirGrid/SONI, CSO, ERCOT and Potomac Economics reports; company disclosures; Cambridge CCAF, IEA, US DOE, Wood Mackenzie, Duke University; figures current to June 2026 unless noted.

Executive Summary

Grids are discarding record volumes of clean energy for want of flexibility, while data-centre demand grows faster than wires can be built. The answer is already running: compute that behaves like a battery.

For two decades data centres have been treated as a grid problem — enormous loads with no agency. This report argues the framing was always wrong, and that the proof has been hiding in plain sight. **Bitcoin mining** has quietly become the most operationally proven flexible load on the planet: a ~20 GW global fleet that dispatches itself against price, curtails in under a minute at zero service cost, and — in ERCOT, the live laboratory — earns audited grid revenue doing so. **Virtual power plants (VPPs)**, the control architecture that turns distributed resources into tradeable power stations, have moved from pilot to tens of gigawatts. And **AI**, the next order of magnitude of demand, is the customer with the deepest pockets — but only part of its workload can behave the way mining does.

The report walks the argument in five steps. First, the gap: Ireland — the sharpest case in the developed world — discarded 14% of its wind in 2024 while data centres took 22% of its electricity, a flexibility failure rather than a wires failure. Second, the machinery: what a VPP actually is, and why 37.5 GW of it already operates in North America. Third, the existence proof: how Bitcoin mining monetises flexibility today, and what the ERCOT receipts actually show — including the gap between price-responsive behaviour (~4.6 GW) and formal registration (~0.24 GW). Fourth, the hard part: AI workloads split into three flexibility classes — training (flexible with engineering), user-facing inference (effectively firm), and batch (flexible by nature) — and a grid-facing control system must treat them separately. Fifth, the synthesis: the **Super-VPP**, a single co-located campus that tiers those workloads behind one grid connection, uses a Bitcoin buffer to make over-built renewables economic, and sells flexibility as a product.

The conclusion is deliberately narrow. Flexible compute is not a theory — the components are all operating somewhere today at commercial scale. What does not yet exist is the full integration, and the economics of that integration depend on an anchor tenant, not on the power markets alone. The prize for getting it right is a data centre the grid pays to host.

THE LENS: FOUR TIERS OF FLEXIBILITY

Firm (user-facing AI inference — never curtailed) · **Semi-flex** (AI training — throttle, checkpoint, migrate) · **Flex** (batch and asynchronous inference — defer for hours) · **Buffer** (Bitcoin mining — first off when the grid is short, first on when renewables spill). Hold these four apart and both the engineering and the economics of flexible compute become legible. The tiers are the report's organising idea, and the Super-VPP in Section 9 is simply the four tiers built deliberately on one substation.

1. The Gap: Clean Energy Discarded Beside Load That Cannot Flex

Start with the problem, stated in one country's numbers. Ireland is the developed world's cleanest illustration because its grid is small, islanded, wind-heavy — and hosts one of the densest data-centre clusters on Earth.

In 2024, **14.0% of all available wind energy across the island was dispatched down** — 2,181 GWh of clean electricity, generated and then discarded (EirGrid/SONI, Annual Renewable Energy Constraint and Curtailment Report 2024). Aurora Energy Research put the cost of that wasted energy to consumers at roughly **€450 million** for the year. And the cause is the crucial detail: **96% of the Republic's wind curtailment** was logged against the system's minimum-conventional-generation limits — the requirement to keep enough synchronous generation running for stability. This is not a shortage of wires. It is a shortage of flexibility: nothing on the demand side is able to absorb the surplus when the wind blows, so the surplus is thrown away.

Beside that waste sits relentless demand growth. **Data centres consumed 22% of Ireland's metered electricity in 2024** (CSO, June 2025) — up from 5% in 2015, more than all urban households combined — and EirGrid's projections put data centres and other new large energy users at roughly **30% by the early 2030s**. The same grid that discards clean energy cannot host new load fast enough; the all-island DS3 system-services budget, capped at €235M a year by the SEM Committee, has in practice been allowed to overrun since 2023 (~€311M in the 2023/24 tariff year) because stability services are that scarce.

Ireland is the sharpest case, but the pattern is global. Renewables dominate new generation everywhere and sharpen net-load ramps; the IEA flags typical transmission lead times of four to eight years, and transformer procurement waits have roughly doubled; PJM — the largest US market — has cleared its capacity auction at the price cap in consecutive auctions. Everywhere the same conclusion is forming: **marginal capacity is now cheaper to aggregate than to build.**

> **THE THREE IRISH NUMBERS**
>
> **14%** of all-island wind dispatched down in 2024 (2.2 TWh, ~€450M of value) · **22%** of national electricity already consumed by data centres · **~30%** by the early 2030s on EirGrid's projections. A grid that discards clean energy while straining to host the load that could absorb it — the gap *is* the opportunity.

2. What a Virtual Power Plant Actually Is

A virtual power plant does three things: it **aggregates** distributed resources — batteries, flexible loads, generation — under one controller; it **dispatches** them in software, in seconds; and it **trades** the result in energy markets as if it were a single power station. The standards layer already exists (OpenADR 3.0, IEEE 2030.5, IEC 61850); the engineering heart is the DERMS — the distributed-energy-resource management system that translates market signals into set-points.

> **A VPP IN PLAIN TERMS**
>
> A VPP is not a building. It is **a control system with a balance sheet**. The distinction that matters is a closed, market-aware control loop — dispatching resources against prices every few seconds — versus an annual interruptibility contract that fires a few times a year. The first is a power station wearing software; the second is a discount for occasionally switching off.

3. VPPs at Scale Today

The common objection — that VPPs are future-tense — is roughly a decade out of date. **North America operates ~37.5 GW of behind-the-meter VPP capacity today** (Wood Mackenzie, September 2025), growing around 14% a year; the US Department of Energy's *Pathways to Commercial Liftoff* projects **80–160 GW by 2030** — 10–20% of peak demand — saving roughly $10 billion a year in avoided infrastructure.

- **Kraken (Octopus Energy Group, UK/EU)** — passed **2 GW** of aggregated residential flexibility in 2025, across 500,000+ consumer devices (EVs, heat pumps, home batteries); one of the largest residential VPPs in the world.
- **AGL (Australia)** — orchestrates ~**1.5 GW** of decentralised assets (FY25), targeting 1.6 GW by 2027, including a 200 MW virtual-battery tolling agreement with Neoen.
- **Sunrun CalReady (California)** — ~75,000 home batteries across 56,000+ households delivering an average of **250 MW per two-hour event** (instantaneous peaks up to 375 MW) — a residential virtual peaker.

Consumer devices, in other words, already work as aggregated grid resources. The open question is what *industrial-scale* asset class belongs on the other end of the DERMS — something large, interruptible, and profitable enough to build merchant. That is where compute enters.

4. Bitcoin: the Largest Battery the Grid Already Has

Seventeen years of Bitcoin have quietly built the largest, most distributed, most flexible electrical load on the planet. The move is to stop treating it as a problem and start treating it as the asset.

The network's vital statistics, as of mid-2026: hashrate around **1 ZH/s** (it first crossed that threshold in September 2025); a continuous global electrical draw of roughly **20 GW** — more than twice Ireland's peak demand; annual consumption of about **138 TWh** (Cambridge Centre for Alternative Finance, April 2025 revised methodology), around half a percent of global electricity; and a measured energy mix of ~**52% renewables plus nuclear**, with natural gas the largest single fuel.

What makes this fleet a de facto VPP is not any single operator's virtue — it is the protocol's economics. The dispatch signal is **price**. Every miner, every block, runs the same calculation: *is my marginal electricity cost below the value of the bitcoin I can mine with it (the "hashprice")?* When power is scarce and expensive, mining is unprofitable and machines switch off; when renewables spill and prices collapse — or go negative — machines switch on. Because the network's difficulty re-targets automatically, the global fleet continuously rebalances toward the cheapest joules anywhere on Earth. The result is **a globally distributed, autonomously dispatched, market-aware load** — the textbook definition of a virtual power plant, baked into the protocol rather than bolted on with a DERMS.

Three engineering properties make mining the gold standard of flexible load, and they are worth stating precisely because Section 7 will measure AI against them. **Response speed:** modern mining firmware pauses and resumes hashing in under a minute — often seconds — with no human in the loop. **Zero service cost:** no customer is waiting for a block to be mined; curtailment breaches no service agreement. **Zero restart cost:** hashing is stateless — there is no work-in-progress to lose, no checkpoint to restore. No other gigawatt-scale load on the grid combines all three.

> **NEGATIVE GENERATION**
>
> From the control room's perspective, 100 MW of mining that switches off in sixty seconds is indistinguishable from a 100 MW power station switching on — but it was built years faster, at a fraction of the capital cost, and it pays its own way between grid events. Grid engineers increasingly call such loads **negative generation**. The reframing matters: it moves the data centre from the demand column to the resource column.

5. The ERCOT Proof: Receipts, Not Theory

Texas is where the thesis stopped being an argument and became an income statement. ERCOT — an energy-only market with scarcity pricing and no capacity payments — is the world's live laboratory for flexible load, and Bitcoin miners were its first industrial students.

The headline receipts belong to **Riot Platforms**, operator of the Rockdale facility (~700 MW) with Corsicana developed behind it. Riot's disclosed power and curtailment credits: **$71 million in 2023** — including **$31.7 million in August 2023 alone**, during the state-wide heat event, a month in which curtailing paid better than mining — and **$33.7 million in FY2024**. These are audited, disclosed figures, not modelled estimates: an existence proof that a data centre can be paid meaningful revenue for flexibility.

Three market mechanisms turn that flexibility into revenue. **Controllable Load Resource (CLR)** registration lets a load qualify as a dispatchable resource in ERCOT's security-constrained economic dispatch — treated, literally, like a generator. **ECRS**, launched June 2023, added a contingency-reserve product that fast-responding loads can sell into. And **four-coincident-peak (4CP) avoidance** rewards loads that dodge the four highest-demand intervals of the year with a year of reduced transmission charges.

> **BEHAVIOUR VS REGISTRATION — THE HONEST NUMBER**
>
> Precision matters here. Potomac Economics' 2025 State of the Market Report puts ERCOT's crypto-mining fleet at **~4.6 GW of aggregate demand**, overwhelmingly price-responsive in behaviour — but only **~240 MW formally registered as CLRs**, with roughly another **1 GW participating through the Emergency Response Service**. The fleet behaves flexibly; formal registration lags well behind. That gap is the point, twice over: the flexibility is real and observable in market data, and the runway for turning it into *registered, contracted* grid capacity remains enormous.

Response times in the field run from seconds to under two minutes — comfortably inside every ancillary product's requirements. And the model has already crossed borders and asset classes, which is where the story turns from mining to compute in general.

Already deployed: the operating fleet

None of this is a pilot programme. The table below is a sample of operating deployments — and, in the middle rows, the pivot that connects Bitcoin's proof to AI's future.

Operator / site	Scale	What it demonstrates
Riot Platforms — Rockdale, TX (ERCOT)	~700 MW	Flexibility as an income statement: $71M (2023) and $33.7M (FY2024) in disclosed power credits.
Crusoe → **Stargate** — Abilene, TX	1.2 GW campus	The pivot in one company: an energy-first operator that began in flared-gas Bitcoin mining now builds and powers AI training at gigawatt scale for OpenAI-class tenants.
Lancium — Abilene Clean Campus, TX	1.2 GW interconnection	Pioneer of ERCOT's load-only controllable-load model (first data-centre CLR, June 2020); now hosts AI tenants on a flexibility-first campus.
IREN — British Columbia	~160 MW	Hydro-powered campuses built for mining, substantially repurposed toward AI/HPC — the site strategy outliving the original workload.
Soluna — Project Dorothy, TX	50 MW (100 MW planned)	Co-located with the 150 MW Briscoe wind farm (acquired outright, April 2026): compute as the buyer for otherwise-curtailed wind.
Gridless — Kenya, Malawi, Zambia	six sites, ~50–500 kW each	The model at the opposite end of the scale: stranded mini-hydro monetised by small modular mining.

Table 1. Operating deployments, mid-2026. Scale figures are approximate and observation-dated; see Sources. The Crusoe, Lancium and IREN rows mark the industry's pivot: energy-advantaged, flexibility-first sites built for Bitcoin now hosting AI — the same architecture serving the customer who pays more.

Note what the middle rows share: the assets built for mining — energy-advantaged land, high-capacity interconnects, flexibility-native operations — turned out to be exactly what AI infrastructure needs. Crusoe kept the architecture and changed the tenant. That pivot is the bridge to the next two sections: how much of AI's workload can actually inherit the flexibility the sites were built around?

6. AI Training: Flexible, With Engineering

To the grid, a frontier training run is not the flat, well-behaved load a data centre is assumed to be. Training alternates between compute and gradient synchronisation on millisecond timescales, so the power draw oscillates violently. Meta's Llama-3-scale runs — roughly 24,000 H100 GPUs, ~30 MW of IT load — exhibited instantaneous swings of tens of megawatts, severe enough that Meta ran *dummy workloads* purely to smooth the electrical profile. At gigawatt scale, modelled oscillations reach hundreds of megawatts, with forcing frequencies in the range that can excite sub-synchronous resonance on the grid (Microsoft/OpenAI/NVIDIA, *Power Stabilization for AI Training Datacenters*, 2025). Utilities have begun writing interconnection requirements around it.

But the same workload has three genuine flexibility levers, all in production today. **Checkpointing:** training state can be saved on a ~10-minute cadence, so a run can pause for a grid event and resume with bounded loss. **Power-capping:** GPUs can be held to 60–80% of thermal design power with a modest throughput penalty — a dimmer switch, not an off switch. **Geographic shifting:** Google's carbon-intelligent computing moves batch workloads between regions and hours in production, and Lancium markets the same capability inside ERCOT. Training is not firm load; it is *schedulable* load with an engineering cost.

> **THE HEADROOM NUMBER**
>
> Duke University's *Rethinking Load Growth* (February 2025) quantified the prize: across the 22 largest US balancing authorities, **~98 GW of new load could be integrated with no new generation or transmission** — if that load curtails just **0.5% of its maximum annual energy consumption**. Most of the required curtailment is partial throttling, not shutdown. For AI operators desperate for interconnection, flexibility is the fastest queue-jump available; for grids, it is nearly-free capacity.

The catch — and it is a serious one — is that "AI" is not one workload. Training's levers do not transfer to the workload that will dominate AI's energy budget by 2030: serving models to users in real time.

7. AI Inference: the Hard Case — and the Trapdoor

User-facing inference is where the flexibility argument gets hard. A chat product answering with sub-100-millisecond service objectives cannot nap because the grid needs frequency response; curtailment is not a throughput penalty but a service breach, visible to customers within seconds. The only levers left are power-capping and graceful degradation — serving slightly slower, or with smaller models — which buys percentage points, not the deep curtailment the buffer tiers deliver. In the four-tier language of this report, user-facing inference is **firm load**, full stop. Any honest flexible-compute architecture must be designed around that fact rather than against it.

The trapdoor is that a large and growing share of "inference" is not user-facing at all. Overnight embedding jobs, evaluation pipelines, batch document processing, asynchronous agents that run for minutes or hours — none of it has a human waiting on the other end. **Batch and asynchronous inference behaves like training:** it can pause at an epoch boundary in seconds, defer for hours, and restart cheaply. Operationally it is flex-tier load wearing inference's name badge.

That is why the report keeps insisting on workload *classes* rather than workload *names*. A DERMS that sees "inference" as one thing will either curtail what it must not, or fail to monetise what it could. The control plane has to distinguish firm, latency-bound serving from batch and asynchronous work — a distinction that lives in the scheduler, not the meter. By 2030 inference in aggregate is projected to dominate AI compute demand, which makes the split decisive: **the largest AI load on the grid in 2030 is currently the least flexible — unless it is re-architected now.**

> **DESIGN RULE**
>
> Never promise the grid flexibility from latency-bound inference. Sell the flexibility of training, batch and buffer tiers; engineer serving capacity as firm load with its own resilience story. Architectures that blur this line fail an audit in both directions — over-promising to the system operator and under-monetising the flexible tiers.

8. The Workload Matrix: Ranking Compute as a Grid Resource

Put the last four sections side by side and the ranking writes itself. Read across the rows: mining is the gold standard on every dimension; training inherits some of its properties at a price; user-facing inference inherits almost none; batch inference quietly inherits nearly all of them.

Property	BTC mining	AI training	User-facing inference	Batch / async inference
Curtailment response	<60 s, firmware	Seconds to throttle; minutes to checkpoint	Power-cap only	Seconds, at an epoch
Service impact	None	Slows time-to-train	**SLA breach**	None within its window
Restart cost	None (stateless)	Progress since last checkpoint	n/a	Low
Geographic flexibility	Fully fungible	High — in production	**Low (latency-pinned)**	High
Cost of curtailment	~$30–80/MWh	~$200–500+/MWh (GPU opportunity cost)	Effectively unbounded	~Training

Table 2. Compute workloads as grid resources, mid-2026. Costs are indicative opportunity-cost ranges, not market quotes.

Two conclusions follow directly. First, **mining's role is structural, not sentimental**: it is the only tier whose curtailment costs the operator almost nothing, which is what makes it the buffer that protects every other tier — and what makes over-built renewables economic, as the next section shows. Second, **the tiers are complements, not competitors.** A campus that runs only mining leaves the anchor revenue on the table; a campus that runs only AI has no cheap flexibility to sell and no buffer to absorb spill. The architecture that pays is the one that stacks them deliberately.

9. From Bitcoin Site to Super-VPP

The growth path starts from sites that already exist — co-located renewables, mining, and an aggregator at sub-100 MW scale — and adds the tenant who pays most: AI. The destination design is a **500 MW co-located campus** on a single 220 kV substation: 300 MWp solar, a 200 MW wind PPA, a 100 MW / 400 MWh battery, a 400 MW import / 200 MW export connection, and the four load tiers stacked behind it.

Tier	Capacity	Behaviour on a grid event
1 · **FIRM** — user-facing inference	150 MW	Never curtailed. The anchor tenant's SLA is untouchable — and the anchor lease is what pays for the campus.
2 · **SEMI-FLEX** — AI training	200 MW	Throttles in seconds; checkpoints and pauses inside two minutes.
3 · **FLEX** — batch / async inference	50 MW	Pauses at an epoch boundary in seconds; defers for hours.
4 · **BUFFER** — Bitcoin mining	100 MW	**First off** when the grid is short; **first on** when renewables spill. Monetises surplus the PPA would otherwise waste.

Table 3. The four-tier load hierarchy. Net grid draw is engineered to roughly 200–300 MW; the campus operates as a two-quadrant grid resource — able to shed load and export power.

On a grid event the campus responds as one machine, on one clock:

```
T + 0-500 ms     Battery delivers full fast-frequency response
T + 0-2 s        Bitcoin buffer drops 100 MW
T + 2-10 s       Batch inference pauses at an epoch
T + 10 s-2 min   Training checkpoints and pauses
T + 2 min        Ancillary contract fulfilled — over 350 MW shed;
                 firm inference untouched throughout
```

The revenue stack (illustrative, annual, power side only): recovered-spill mining ~€45M · capacity value on the sheddable 350 MW ~€38M · avoided-curtailment value to the PPA partner ~€25M · fast frequency response from the battery ~€15M · wholesale arbitrage ~€7M — **roughly €130M a year**, before the anchor lease, which is an order of magnitude larger. In Satscryption's techno-economic model (v1.6, Ireland base case) the power stack alone is near break-even — ~7.85% unlevered project IRR, ~18.9% levered equity IRR with the tenant-funded shell — and the streams earn from each other: **without the Bitcoin buffer the over-built renewable PPA is uneconomic; without the anchor tenant there is no engine; without the battery, training's oscillations damage the grid.** One integrated campus beats a fleet of separately optimised assets.

10. Where It Works, and Where It Does Not

Strong fit	Poor or no fit
Grids with high renewable curtailment and system-stability limits (Ireland, ERCOT, South Australia)	Grids with abundant firm capacity and no scarcity pricing
Workloads that checkpoint, defer or relocate — mining, training, batch inference	Latency-bound, user-facing inference promised as "flexible"
Sites with over-built renewables seeking a buyer for spill	Standalone data centres with no generation or storage partner
Markets with controllable-load and ancillary programmes (the ERCOT template)	Markets offering only annual interruptibility contracts

Table 4. The decision boundary. Economic deterrence of curtailment is a layer, not a universal answer: the model pays where flexibility is scarce and priced.

11. Risks and Honest Caveats

- **Design-stage economics.** The campus numbers are modelled, not operated: in the v1.6 base case the power stack alone is near break-even (project NPV modestly negative at WACC) and returns hinge on the AI anchor lease. **Tenant risk is the dominant risk.**
- **Training's electrical signature.** Gigawatt-scale oscillation ("AI flicker") is a real interconnection concern; the battery and power-stabilisation engineering are load-bearing parts of the design, not options.
- **Inference must be engineered around.** The firm tier can never be sold as flexibility; architectures that blur the classes fail in both directions.
- **Bitcoin-side variance.** Network statistics carry estimation uncertainty, and buffer-tier revenue moves with hashprice; the buffer's value is real but volatile.
- **Regulatory openness.** Capacity-market treatment (de-rating) of flexible compute remains unsettled in most markets; ERCOT's programmes are a template, not a universal fact.

12. Where This Stands Today

Every component of the flexible-compute thesis is operating somewhere at commercial scale in mid-2026. Consumer VPPs run at tens of gigawatts. Bitcoin mining earns audited grid revenue at sub-minute response in ERCOT, with ~4.6 GW behaving flexibly against ~1.2 GW formally enrolled across CLR and ERS. Training flexibility — checkpointing, power-capping, geographic shifting — is in production at Google, Meta-scale labs and Lancium. Batch inference defers by the hour as a matter of routine. The 500 MW integrated campus, tiering all of it behind one substation, is the one piece that exists only as a validated design — its components proven separately, its integration not yet operated.

> **THE MODEL — USE IT, BUT KNOW ITS EDGES**
>
> The Super-VPP economics in this report come from Satscryption's v1.6 techno-economic model: a 15-year, tier-aware co-optimisation of the campus's energy dispatch and revenue stack, adversarially reviewed and cross-checked against a shadow model. It remains a model — its base case assumes an anchor tenant on a powered-shell lease, Ireland market prices, and mid-2026 capital costs. Treat the structure of its conclusion (power stack ≈ break-even; anchor lease = the engine; buffer = the PPA's economics) as the signal, and the decimals as inputs to re-run, not facts to quote.

Bottom line. For two decades, data centres have been the grid's problem child — the load that only ever asked for more. Bitcoin mining, almost by accident, demonstrated the alternative: compute that behaves like a battery, absorbing surplus and vanishing in scarcity, paid for the service. AI raises the stakes in both directions — its training and batch tiers are the largest new flexible load available to any grid on Earth, while its serving tier is as firm as load gets. The Super-VPP is simply the architecture that takes all of that seriously at once: four tiers, one substation, flexibility sold as a product rather than suffered as a constraint. **Bitcoin built the rails. AI is the customer who pays to scale them.** The grids that price flexibility — and the operators who productise it — will find that the largest battery on the grid was already plugged in.

Sources, Methodology & Confidence

This report synthesises regulator and system-operator publications, company disclosures, and independent research, prioritised in that order. Figures are point-in-time — most current to June 2026 — and observation-dated; grid and market statistics move continuously. All claims were adversarially fact-checked against primary sources in July 2026; where a figure is an estimate or a modelled result, it is flagged in the text.

Grid & market — primary

• EirGrid/SONI, *Annual Renewable Energy Constraint and Curtailment Report 2024* (April 2025) — 14.0% all-island wind dispatch-down (2,181 GWh of 15,570 GWh available); 96.4% of Republic wind curtailment under the High-Frequency/Minimum-generation reason code. €450M consumer-cost estimate: Aurora Energy Research (reported July 2025).

• CSO Ireland, *Data Centres Metered Electricity Consumption 2024* (June 2025) — 22% of national metered electricity (5% in 2015). EirGrid capacity outlooks — data centres and new large users ~28% by 2031, ~31% by 2034 (median).

• SEM Committee — DS3 system-services expenditure cap €235M/yr (SEM-17-080); overruns permitted since 2023, ~€311M in 2023/24.

• Potomac Economics, *2025 State of the Market Report for ERCOT* — crypto-mining demand ~4.6 GW; ~240 MW registered CLR; ~1.05 GW via ERS. PJM Base Residual Auction reports — consecutive cap-clearing auctions ($329.17 then $333.44/MW-day UCAP).

Companies & deployments

• Riot Platforms disclosures — $71M (2023), $31.7M (Aug 2023), $33.7M (FY2024) in power and curtailment credits.

• Crusoe / Stargate Abilene (1.2 GW); Lancium (first data-centre CLR, June 2020; 1.2 GW approved interconnection); IREN British Columbia (~160 MW, repurposing toward AI); Soluna Project Dorothy (50 MW; 150 MW Briscoe wind farm acquired April 2026); Gridless (six sites, Kenya/Malawi/Zambia). Company releases and site disclosures.

• Kraken/Octopus (2 GW residential flexibility, 2025); AGL ESG Data Centre FY25 (1,487 MW under orchestration); Sunrun CalReady (May 2025: ~75,000 batteries, 250 MW average / 375 MW peak).

Research & independent analysis

• Cambridge Centre for Alternative Finance (April 2025, revised methodology) — ~138 TWh/yr; ~52% renewables + nuclear.

• IEA, *Energy and AI* (April 2025) — 415 TWh (2024) → ~945 TWh (2030); US + China ~80% of growth. AI share of data-centre power 5–15% → 35–50%: Kamiya & Coroamă (EDNA / IEA 4E TCP, March 2025).

• Norris et al., *Rethinking Load Growth* (Duke Nicholas Institute, February 2025) — 98 GW at a 0.5% average annual curtailment rate (share of the new load's maximum annual energy).

• Meta engineering disclosures (Llama-3-scale training); SemiAnalysis, *AI Training Load Fluctuations at Gigawatt-scale* (June 2025); Microsoft/OpenAI/NVIDIA, *Power Stabilization for AI Training Datacenters* (arXiv:2508.14318).

• US DOE, *Pathways to Commercial Liftoff: Virtual Power Plants*; Wood Mackenzie VPP market monitor (September 2025).

• Satscryption, *VPP Financial Model v1.6* (June 2026) — 500 MW campus techno-economics; base case 7.85% project IRR / 18.9% equity IRR, tenant-funded shell, Ireland region. Modelled, not operated.

• Satscryption keynote film, *The Largest Battery the Grid Already Has* (8 June 2026) — youtu.be/ue8tEJeHCWk; companion Satsight research note SSC-RN-2026-013.

Disclaimer. This document is an informational research synthesis prepared for general understanding. It is not investment, legal, engineering, or financial advice, and not a recommendation regarding any security, digital asset, or infrastructure project. Modelled results (including the v1.6 campus economics) are analytical projections, not measured outcomes. All figures are dated and should be verified against primary sources before any decision; market, regulatory and tax details are summarised and approximate.