

Can Flexible Compute Carry Debt? AI Workloads as an Electricity Asset

A financing framework for converting controllable computing demand into bankable electricity value

Authored by Chennakeshav Adya

Independent Researcher

October 2026

ABSTRACT

AI data centres are usually analysed as large electricity consumers. Some workloads can also be scheduled, slowed, relocated or interrupted within defined service constraints. That controllability can reduce coincident peak demand, support grid operations, improve the use of constrained connections and defer selected network investment. The central financing question is whether those benefits create cash flow that a lender can recognise. This paper develops a Flexible Compute Bankability Framework for data-centre sponsors, utilities, grid operators, infrastructure investors and lenders. It separates physical capability from dispatch authority, measured performance from theoretical flexibility, and contracted value from merchant optionality. The framework evaluates six linked elements: workload eligibility, telemetry and control, customer service obligations, electricity-market access, contractual payment design and lender protections. Official studies and field demonstrations show that data-centre demand can respond to grid conditions, but the available flexibility varies by workload, facility design and operating agreement. Three author-developed figures and three scenario tables illustrate how a controllable compute portfolio might translate verified megawatts into operating savings or contracted availability payments. The scenarios are hypothetical. They do not represent observed project economics, a forecast, a credit opinion or an investment recommendation.

JEL Classification: G21, G28, G32, L86, L94, Q41, Q48

Keywords: artificial intelligence, data centres, flexible compute, demand response, large loads, project finance, infrastructure debt, electricity markets, grid interconnection, capacity payments

Introduction

Electricity has become a gating factor for AI infrastructure. A developer can secure land, permits, fibre and computing equipment while waiting years for sufficient grid capacity. Utilities face equally difficult choices. They must plan generation, transmission and distribution for loads that can be exceptionally large, can ramp quickly and can change before the supporting infrastructure is recovered through tariffs. Regulators must protect existing customers from stranded costs without deterring investment.

Flexible compute offers a potential bridge. Model training, batch inference, data processing, checkpointing and selected development workloads may tolerate changes in time, place or power level. Cooling systems, batteries and onsite generation can add further controllability. The combined facility can reduce demand during a system event, follow a negotiated schedule or absorb electricity when network capacity is available.

Technical flexibility does not automatically create financeable value. A lender requires a defined obligor, payment formula, dispatch protocol, performance test, liability regime and evidence that curtailment does not destroy the revenue used to service debt. The financing case also depends on who controls the workload. A building owner cannot promise the flexibility of a tenant whose service-level agreement prohibits interruption. A cloud operator cannot count a training job twice, first as customer capacity and again as committed grid response.

This paper treats flexible compute as an electricity asset only when its physical and contractual attributes can be measured, controlled and enforced. It asks how that asset should be diligenced, contracted, valued and incorporated into debt sizing.

1 Define the asset precisely

The relevant asset is a verified right to alter net electricity demand within agreed limits. It is not the server, the software or the grid connection in isolation. Its value arises from the coordinated ability to change power consumption at a specified location and time while preserving agreed computing outcomes.

That right contains several dimensions. Capacity states how many megawatts can move. Duration states how long the change can last. Response time states how quickly the facility can act. Recovery describes the rebound when deferred computing resumes. Availability defines the hours in which the capability can be called. Location

determines which network constraint or market product it can address. Performance measurement determines whether delivery can be proven.

A credit committee should reject a single headline such as “50 MW flexible”. It should require a dispatch envelope. A facility might provide 50 MW for fifteen minutes, 20 MW for two hours and 8 MW for six hours. Each product has a different operational effect and potential payment. The dispatch envelope should also show exclusions, including customer deadlines, maintenance windows, equipment limits, thermal constraints and minimum stable loads.

2 Separate workloads by operational tolerance

AI workloads are not uniform. Interactive inference can be latency-sensitive. Safety, healthcare, payment and communications applications can require high continuity. Training jobs may tolerate checkpointing, but interruption can waste completed computation or extend delivery dates. Research and development queues can be more flexible. Data preparation, synthetic-data generation and selected batch processes may move across hours or regions when inputs and security policies permit.

The sponsor should classify each workload against four controls: time shift, geographic shift, power modulation and interruption. The classification requires evidence from the workload owner. A data-centre landlord cannot infer flexibility from server type. A colocation tenant may use identical accelerators for services with entirely different tolerances.

The classification must also identify the recovery obligation. Moving a job out of a peak hour can cause a higher load later. Grid value depends on whether that rebound occurs outside the constrained period. Financeable flexibility therefore requires an energy budget as well as a power reduction.

Table 1. Workload classification for a financeable flexibility envelope

Workload category	Potential control	Evidence required	Principal constraint	Financing treatment
Interactive critical inference	Limited power capping or redundant-region transfer	Service-level terms, latency tests, redundancy evidence	Customer continuity and latency	Exclude unless specifically contracted
Batch inference and analytics	Time shift, queue management, regional transfer	Deadline tolerance, data-location permissions, measured completion impact	Delivery window and data governance	Recognise only tested recurring capacity
Model training	Checkpoint, pause, slow or relocate	Checkpoint integrity, restart time, accelerator utilisation and customer consent	Lost work and schedule extension	Apply performance and availability haircuts
Development and testing	Queue, pause or reschedule	Workload inventory and operating policy	User productivity	Potentially high flexibility; low revenue certainty
Cooling and auxiliary systems	Thermal storage, temperature reset and equipment sequencing	Engineering study, environmental limits and commissioning tests	Equipment reliability and thermal rebound	Recognise separately from IT load

3 Build the power-to-compute control chain


The control chain begins with a grid or utility signal. A facility controller translates that signal into an authorised response. The scheduler selects eligible workloads. Server, cooling, battery and generation controls execute the response. Revenue metering and operating telemetry prove performance. A settlement process converts performance into payment or tariff savings.

Every handoff can fail. The market signal can arrive late. The scheduler can select an ineligible job. A battery can be unavailable. Meter data can be incomplete. A customer can object to a delayed workload. The bankability review should map each failure to an owner and remedy.

Cybersecurity is integral. A control path that can reduce a data centre's load can affect computing delivery and critical electrical equipment. Access rights, segmentation, change control, fallback modes, incident response and audit logging should meet an appropriate operational-technology security standard. The dispatch interface should

use the minimum authority required and should preserve safe local control.

Figure 1. Power-to-compute architecture for a controlled flexibility product



This author-developed diagram shows a general control and evidence chain. Actual architectures depend on market, facility and workload design.

4 Distinguish four sources of economic value

Flexible compute can create value through four channels. The first is an availability payment for standing ready to reduce or shift demand. The second is a performance payment when the facility responds. The third is an operating saving, such as avoiding a peak-demand charge or buying electricity in a lower-cost interval. The fourth is connection value, where agreed flexibility enables earlier or larger access to a constrained network.

These channels have different credit quality. A fixed availability payment from a creditworthy utility under a long-term agreement can resemble contracted infrastructure revenue. A market performance payment depends on event frequency, clearing prices and delivered volume. An operating saving reduces cost but does not create a separate receivable. Connection value may be commercially important while remaining difficult to monetise as annual cash flow.

The cash-flow model should prevent double counting. A facility should not treat avoided demand charges, wholesale-market revenue and a utility payment as independent if one dispatch event drives all three and programme rules prohibit stacking. It should also deduct the economic cost of delayed computing, additional energy use, battery degradation, maintenance, control software and market participation.

5 Define a credible baseline

Performance is measured against what the facility would have consumed without dispatch. That counterfactual is a major source of model risk. Data-centre load can change with customer activity, model stage, cooling conditions, equipment availability and commissioning. A static baseline can overstate or understate delivered flexibility.

The measurement plan should specify the baseline methodology before financial close. It should state the meter boundary, interval length, adjustment for weather or production, treatment of outages, data retention and dispute process. Historical simulations should show how the method behaves during normal and unusual periods. Independent metering or an agreed market settlement system strengthens evidential quality.

Lenders should focus on false positives. If a workload would have ended naturally during an event, its decline should not be paid as response. If the operator increases load before the baseline window, the programme can create an incentive to inflate the counterfactual. Governance should prohibit baseline manipulation and give the offtaker or programme administrator audit rights.

6 Verify response through staged commissioning

Flexibility should be commissioned as rigorously as generation or storage. Testing should begin with component controls, then workload groups, then the entire facility. Each test should measure response time, delivered megawatts, duration, recovery, computing impact and equipment alarms.

A single demonstration is insufficient for debt sizing. The facility should repeat tests across representative operating states. High utilisation, low utilisation, hot weather, cooling maintenance and partial equipment outages can produce different results. The test plan should also include failed communication, rejected dispatch and emergency restoration.

Acceptance should produce a certified capacity schedule rather than a binary pass. The schedule can state 20 MW available continuously, another 15 MW subject to workload conditions and 10 MW available only through storage. Lenders can apply different recognition factors to each layer.

7 Allocate dispatch authority

The contract must identify who can call an event, who can refuse and who bears the consequence. Grid emergencies may require rapid action. Economic dispatch can allow more discretion. A customer-service event may override electricity optimisation. Ambiguous priority rules create both operational and credit risk.

The facility should use a hierarchy. Safety and legal compliance come first. Critical customer obligations follow. Firm grid commitments then rank ahead of voluntary market optimisation. The hierarchy should be coded into control systems and reflected in customer, utility and financing documents.

Dispatch authority should also be limited. The utility should not receive unrestricted control of servers. It should issue a power target, duration and notice. The operator should choose the safe method of delivery. Exceptions should be logged and settled under predefined rules.

8 Align customer contracts

Customer contracts determine whether compute can move. The operator should review service levels, completion deadlines, data residency, security, reserved capacity, power commitments and remedies. A grid agreement cannot override a tenant's rights unless the tenant contract permits the response.

There are several alignment models. A customer can buy a lower-cost flexible service with defined delay tolerance. It can share electricity savings or programme revenue. It can authorise specified workload classes while preserving a firm tier. It can also pay a premium for non-curtable capacity, leaving other capacity available for response.

The operator must avoid selling the same capacity as both firm compute and firm electricity flexibility. Capacity reservation, utilisation reporting and dispatch logs should reconcile the two products. Customer consent should be explicit where performance can be affected.

9 Structure the utility or grid contract

A bankable contract needs a clear product definition. The agreement should specify contracted megawatts, notice, ramp rate, duration, annual event cap, recovery window, availability hours, permitted exclusions and testing. Payment should distinguish availability from performance and state how non-delivery is calculated.

Term matters. A short pilot can prove capability but provides limited debt support. A ten-year agreement can support financing if the payment obligation, regulatory treatment and termination regime are credible. The lender should test whether the utility can recover programme costs and whether a regulator can alter the tariff.

Termination compensation should reflect remaining debt exposure where the offtaker ends the contract without operator default. Change-in-law and market-rule provisions should address a product becoming unavailable. Assignment and lender step-in rights should be agreed before close.


10 Evaluate connection value separately

Flexible interconnection can allow a large load to connect sooner or at a higher nominal capacity if it accepts curtailment under defined network conditions. The value can be substantial because delayed energisation can defer computing revenue and strand equipment commitments.

The connection agreement should state the firm import level, conditional import level, constraint that triggers curtailment, expected event profile and upgrade path. It should identify whether curtailment is compensated. An earlier connection with uncompensated interruption can still create value, but that value appears in project timing and utilisation rather than programme revenue.

Debt sizing should use the firm connection as the base case unless the conditional capacity has an enforceable operating history or a contractual minimum availability. The model should stress a year with materially more curtailment than expected. It should also reflect the risk that network upgrades arrive later than planned.

Figure 2. Illustrative connection and bankability timeline



This author-developed timeline is hypothetical. Dates and milestones are illustrative and do not describe a specific project.

11 Translate value into lender-recognised cash flow

Lenders should begin with gross contracted value and apply deductions. Availability revenue should be reduced for expected outages, workload ineligibility, performance failures, operating cost, customer compensation and taxes. Merchant revenue should receive a lower recognition factor or remain outside the debt-sizing case.

Operating savings require careful treatment. A lower electricity bill can improve debt service coverage, but only if the baseline tariff, dispatch behaviour and production effect are credible. The lender should compare savings with the lost gross profit from delayed or reduced computing. If the project earns revenue by accelerator-hour, interruption can reduce revenue even while saving power cost.

The recognised amount should be limited to the lower of contractual entitlement, demonstrated performance and available workload flexibility. A 30 MW contract supported by 18 MW of repeatable response should be modelled at the evidence-based level until further capacity is accepted.

Table 2. Illustrative annual value bridge for a flexible-compute programme

Item	Hypothetical base case	Lender question	Illustrative treatment
Contracted availability	USD 5.4m	Is the payment fixed, creditworthy and enforceable?	Start with contracted amount
Expected performance payments	USD 1.8m	How often are events called and settled?	Apply a conservative recognition factor
Peak and tariff savings	USD 2.2m	Are savings measured after rebound and production loss?	Recognise tested recurring savings
Customer compensation and lost contribution	USD (1.7)m	Does curtailment reduce compute revenue or require credits?	Deduct in full in base case
Control, metering and maintenance cost	USD (0.9)m	Are all incremental operating costs included?	Deduct in full
Merchant optimisation	USD 1.5m	Is the product exposed to price and rule changes?	Treat as upside or heavily haircut


12 Determine debt capacity

Additional debt capacity depends on recognised net cash flow, required debt-service coverage, interest rate, amortisation and contract life. A long-lived asset with a short flexibility agreement should not capitalise payments beyond the contracted term without a robust renewal assumption.

The debt model should use the shorter of the payment contract and the reliable operating life of the required controls or equipment. It should include replacement expenditure for batteries, switchgear, control systems and metering where those assets enable delivery. Residual value should exclude uncontracted flexibility unless supported by an observable market.

Cash flow can support project debt directly, reduce the reserve required for electricity-price exposure or improve covenant headroom. The financing benefit may therefore appear in structure rather than a separate debt tranche. Lenders should document the mechanism to avoid overstating leverage.

Figure 3. Illustrative debt capacity sensitivity to verified flexible load and net annual value



This author-developed sensitivity uses hypothetical assumptions: ten-year amortisation, 8.0 per cent interest and 1.40 times debt-service coverage. It is not a valuation or financing offer.

13 Apply recognition tiers

A practical financing framework can use three tiers. Tier one includes fixed availability payments under an enforceable contract with a creditworthy counterparty, supported by accepted performance tests. Tier two includes demonstrable tariff savings and expected performance payments with sufficient operating history. Tier three includes merchant optimisation, avoided future upgrades and uncontracted connection value.

Tier one can enter the base debt case after operating costs and performance deductions. Tier two can receive a haircut or support covenant headroom. Tier three should usually remain sponsor upside until it becomes contracted or evidenced through a stable market.

The tiers should be reviewed periodically. Successful dispatch history can move value upward. Customer concentration, programme rule changes or declining workload flexibility can move it downward. Financing documents should require reporting that supports this reassessment.

14 Design performance security

The operator may owe damages for non-delivery. Unlimited exposure can undermine project credit. The contract should use proportional performance deductions, defined caps and force-majeure treatment consistent with the operator's control.

Security can include a letter of credit, reserve account or parent guarantee. Its size should reflect expected exposure rather than the gross value of the facility. The lender should ensure that utility security does not rank ahead of essential project obligations without limit.

Performance security should step down after a reliable operating history where appropriate. It should also distinguish a failed market response from a computing outage. The same event should not generate duplicative claims beyond agreed caps.

15 Integrate flexibility into the financing documents

The common-terms agreement should recognise the flexibility contract as a material project document when it supports debt. Amendments, waivers and termination should require lender consent. The security package should include assignment of receivables and rights under the agreement, subject to public-law limitations.

Financial covenants should identify the recognised cash-flow amount. Distribution tests can exclude merchant flexibility revenue or require it to be trapped until performance thresholds are met. Reserve accounts can cover delayed settlement or seasonal event risk.

Information undertakings should include dispatch events, availability, settlement, customer impacts, control-system incidents and programme-rule changes. A lender technical adviser should be able to inspect test records and metering data.

16 Address technology and obsolescence

AI hardware and software evolve quickly. New accelerators can draw more power per rack while delivering more computation per unit of energy. Training methods can change checkpoint behaviour. Workload orchestration can improve geographic and temporal flexibility. These changes can increase or reduce the contracted capability.

The project should maintain a change-control process. New hardware should be tested before its load is included in the flexibility commitment. Software updates to schedulers and energy controls should preserve auditability and fallback modes. Capacity should be recertified after material changes.

Debt should not assume that flexibility rises automatically with computing demand. A larger load can contain a smaller flexible share if it serves more critical inference. The relevant metric is verified controllable megawatts under current contracts.

17 Account for cooling and thermal inertia

Cooling can contribute to response through chilled-water storage, temperature set-point changes, pump sequencing and thermal inertia. These measures can reduce facility power without interrupting computing. Their duration depends on climate, load, equipment and environmental limits.

Thermal response must remain within manufacturer specifications and customer requirements. A short reduction can create a later rebound. The energy model should capture both. Repeated temperature cycling can affect equipment and should be assessed through engineering evidence.

Cooling flexibility should be measured separately from IT flexibility. This prevents a claim based on the same megawatts being counted twice and allows lenders to see which layer remains available when workloads are inflexible.

18 Treat storage and generation as distinct assets

Batteries can reduce grid import while computing continues. Backup generation can also support net-load reduction where permits, fuel, emissions and operating rules allow. These resources may provide more predictable response than workload control, but they have different costs and risks.

A battery has energy limits, degradation and replacement needs. Backup generators can face restrictions on non-emergency operation. Fuel availability and emissions compliance can constrain duration. The financing model should allocate capital and operating cost to the product each asset supports.

The project should report gross workload reduction and net meter response separately. This allows the utility and lender to understand whether performance came from compute, cooling, storage or generation.

19 Evaluate geographic load shifting

Operators with multiple regions can move eligible workloads from a constrained grid to another location. This can create more flexibility than a single facility, but it introduces network, data, latency and carbon considerations.

The receiving region must have computing, electricity and network capacity. Data transfer can consume time and energy. Customer contracts and regulation can restrict data location. A simultaneous grid event across regions

can remove the assumed diversity.

A financing case should use committed reciprocal capacity or demonstrated orchestration rather than theoretical fleet capacity. Intercompany arrangements should define cost allocation and availability. Lenders to one facility need confidence that another affiliate cannot withdraw support during a stress period.

20 Examine market access

Participation rules differ by jurisdiction. Some markets allow aggregated demand response. Others use utility tariffs, emergency programmes or bilateral agreements. Minimum size, telemetry, registration, testing and settlement requirements can determine whether a technically flexible site can earn revenue.

The sponsor should prepare a market-access memorandum identifying eligible products, regulatory approvals, aggregation rights and rule-change exposure. It should confirm whether the same resource can receive multiple payments. The review should also identify taxes, levies and licensing requirements.

Where direct market access is uncertain, a bilateral utility agreement can provide a clearer path. The agreement should still align with the system operator's rules and regulatory cost recovery.

21 Protect other electricity customers

Large-load contracts receive public scrutiny because network investment can outlast the customer. Minimum bills, collateral, long contract terms and exit charges can protect other ratepayers. Flexible-load discounts should reflect demonstrable system value rather than transfer cost to other users.

The facility should quantify avoided or deferred system cost using the utility's planning methodology. The analysis should distinguish local distribution, transmission and generation effects. A reduction during a few annual peak hours can be valuable, but it does not remove the need for firm capacity during all other hours.

Transparent performance reporting can support regulatory approval. The contract should allow adjustment when the promised value is not delivered, while preserving financeability through objective rules.

22 Model downside cases

The credit case should stress physical, commercial and regulatory failure. Physical cases include lower flexible capacity, slower response, cooling constraints and battery degradation. Commercial cases include fewer eligible workloads, customer opt-out, lower market prices and increased compensation. Regulatory cases include programme redesign, baseline changes and higher collateral.

The most important combined case is high electricity cost with low flexibility. This can occur when workloads become critical during a constrained period. The project then loses programme revenue while paying peak prices. Debt service should remain protected under a defined severe case.

Table 3. Downside matrix and financing response

Downside	Transmission to cash flow	Early indicator	Financing response
Eligible workload falls by 40%	Lower availability and performance revenue	Rising critical-inference share	Capacity recertification and cash-flow haircut
Event hours double	More customer disruption and rebound energy	Utility system alerts and constraint notices	Event cap, compensation formula and operating reserve
Baseline methodology changes	Lower settled megawatts	Market-rule consultation	Change-in-law protection and covenant rebase
Customer rejects dispatch	Lost flexibility or service credits	SLA complaints and opt-out requests	Explicit customer election and revenue sharing
Control system unavailable	Non-delivery and potential penalty	Failed tests or telemetry gaps	Redundancy, cyber controls and performance reserve

Downside	Transmission to cash flow	Early indicator	Financing response
Market price declines	Lower merchant revenue	Auction and settlement trends	Exclude merchant value from base debt case

23 Conduct lender technical diligence

The lender technical adviser should review electrical one-lines, metering, control architecture, workload inventory, cooling capability, storage, generation and communications. It should witness representative performance tests and reconcile the result with contract commitments.

The adviser should test the dispatch envelope under equipment outages and high utilisation. It should review recovery energy and workload completion. It should verify that the proposed response does not conflict with equipment warranties, permits or safety requirements.

The report should state certified capacity by duration and notice period, not a single undifferentiated figure. It should list open conditions and a recertification schedule.

24 Conduct commercial and legal diligence

Commercial diligence should examine customer concentration, workload mix, pricing, service credits and renewal. It should identify who benefits from flexibility and whether customers have an incentive to participate. The forecast should reconcile compute revenue with electricity programme revenue event by event.

Legal diligence should review utility authority, regulatory approval, programme rules, customer consent, data governance, cybersecurity obligations, assignment, termination and liability. The review should identify whether payments can be changed through a tariff proceeding.

Where public bodies are counterparties, counsel should confirm appropriation, sovereign-immunity and enforcement considerations. The financing should not assume remedies that are unavailable under applicable law.

25 Establish reporting and governance

A flexibility committee can bring together data-centre operations, workload engineering, customer service, electricity procurement, finance, legal and cybersecurity. It should approve capacity commitments, review performance and govern changes.

Monthly reporting should show eligible megawatts, tested megawatts, event availability, delivered response, settlement, customer impact, rebound energy and incidents. Quarterly reporting should reconcile recognised cash flow with the financing model. Material deviations should trigger recertification or a covenant review.

The board should receive a concise view of contracted value, delivery risk and customer exposure. A strong dashboard avoids treating technical potential as earned revenue.

26 Use a phased financing approach

A greenfield project may have no operating history. It can still finance the underlying data centre on conventional contracted revenues, while treating flexibility as upside. After commissioning, a performance reserve or accordion can release additional debt when tests and settlement history meet defined conditions.

The phases can be: design qualification, control commissioning, initial programme acceptance, twelve months of operations and mature portfolio optimisation. Each phase has an evidence gate. The debt instrument can specify the information and ratios required for an increase.

This approach aligns capital with demonstrated capability. It also preserves an incentive for the sponsor to complete customer alignment and operating controls.

27 Consider portfolio aggregation

A portfolio can aggregate smaller flexible loads, diversify workload patterns and offer a more reliable product. Aggregation can also reduce transaction costs and provide geographic coverage. The benefit depends on correlation. Facilities exposed to the same weather, market and customer demand can fail together.

The aggregator should disclose resource composition, dispatch rights, baseline methods and concentration. It should have back-to-back contracts and avoid promising more capacity than the portfolio can deliver. Settlement and cash allocation should be transparent.

Portfolio debt can recognise diversification after sufficient operating evidence. Lenders should retain site-level reporting and concentration limits.

28 Build a decision framework

The investment committee should ask six questions. Is the workload demonstrably controllable? Does the operator hold the necessary dispatch rights? Can performance be measured against a credible baseline? Is value paid under an enforceable arrangement? Does curtailment preserve the computing revenue that supports debt? Can the project survive a severe combined downside?

An affirmative answer to all six can support lender recognition. A weakness in one area does not always end the project, but it changes treatment. Uncertain market prices can be excluded from base cash flow. Limited history can be addressed through staged drawdown. Missing customer consent requires contractual remediation before commitment.

The decision should record the recognised megawatts, value per megawatt, haircut, term and recertification conditions. This creates an auditable bridge from engineering evidence to debt sizing.

29 Recognise common failure modes

The first failure is using installed load as flexible capacity. The second is relying on software capability without customer rights. The third is measuring response against a weak baseline. The fourth is ignoring rebound energy and delayed computing. The fifth is capitalising merchant revenue as though it were contracted.

Other failures include double counting batteries and workloads, overlooking cybersecurity, assuming regulatory approval, underestimating customer compensation and using a long asset life for a short programme. Each failure can turn a persuasive technical story into a weak credit.

The remedy is evidence discipline. Every recognised dollar should trace to a controlled asset, valid right, measured performance and enforceable payment or saving.

30 Apply an execution sequence

The sponsor should begin with a workload and contract inventory. It should map the electrical and control architecture, define dispatch products and engage the utility or market operator. Customer terms should be aligned before capacity is offered.

The project should then implement telemetry, controls and cybersecurity; conduct staged tests; negotiate settlement and liability; and prepare an independent certification package. The financing model should recognise value only after these steps establish a credible evidence chain.

Following close, governance should maintain the link between technology, customers and finance. New workloads, hardware and market rules should trigger review. Flexibility is an operating capability that requires continuous evidence.

31 Limitations

This paper provides a general corporate-finance and project-finance framework. Electricity markets, tariffs, interconnection rules and customer rights differ by jurisdiction. Market products and regulatory positions can change. Project-specific engineering, legal, tax, accounting, cybersecurity and regulatory advice is required.

The financial figures in the tables and figures are hypothetical author-developed scenarios. They illustrate method. They are not observed data, a forecast, a credit rating, a valuation or a financing offer.

Conclusion

Flexible compute can behave like an electricity asset when it has a verified dispatch envelope, enforceable rights, credible metering and a payment or saving that survives operating costs and computing impacts. The bankable unit is not theoretical megawatts. It is contracted and demonstrated net-load response within a defined time, location and service boundary.

The strongest financing structures separate fixed availability, performance revenue, operating savings, connection value and merchant upside. They align utility, customer and lender documents. They test workload and facility controls under realistic conditions. They also preserve conservative debt service when flexibility is unavailable.

For sponsors, the opportunity is broader than programme revenue. A well-designed flexibility proposition can support earlier connection, improve electricity procurement, strengthen utility relationships and create a differentiated computing product. For lenders, recognition should follow evidence. Flexible compute can carry debt to the extent that control, contract and cash flow remain joined.

Frequently Asked Questions

What does it mean to treat flexible compute as an electricity asset?

It means treating a verified right to change net electricity demand as a measurable service. The service requires defined capacity, duration, response time, availability, location, control and settlement.

Can all AI workloads be curtailed?

No. Flexibility depends on workload deadlines, latency, checkpointing, data location, customer terms and operational risk. Each workload class requires evidence and explicit authority.

Which flexibility revenue is most suitable for debt sizing?

Fixed availability payments under an enforceable long-term agreement with a creditworthy counterparty generally provide stronger support than merchant optimisation. Recognition still requires tested capacity and deductions for operating and performance risk.

Can electricity savings support debt even without a separate payment?

Tested recurring savings can improve debt-service coverage. The model should deduct rebound energy, lost computing contribution, customer compensation and the cost of controls or storage.

How should a lender treat a flexible grid connection?

The firm import level should anchor the base case. Conditional capacity can support upside or additional debt after contractual availability and operating history demonstrate its reliability.

Why is the baseline important?

Performance is the difference between actual demand and the demand that would have occurred without dispatch. A weak baseline can misstate delivery and settlement.

What evidence should the lender require?

The lender should require workload rights, technical design, representative tests, interval telemetry, settlement history, customer alignment, cyber controls and a contract that defines payment and non-delivery.

What is the principal financing risk?

The principal risk is a mismatch between electricity commitments and computing obligations. The project can lose flexibility revenue and compute revenue if dispatch rights, customer terms and operating controls are not aligned.

References

1. Lawrence Berkeley National Laboratory, DOE Data Center Load Flexibility Workshop Summary, January 2025. <https://eta.lbl.gov/publications/doe-data-center-load-flexibility>
2. Lawrence Berkeley National Laboratory, Demand Response Opportunities and Enabling Technologies for Data Centers. <https://eta.lbl.gov/publications/demand-response-opportunities-and>
3. Lawrence Berkeley National Laboratory, Data Centers Research Area. <https://eta-publications.lbl.gov/research-areas/data-centers>
4. Lawrence Berkeley National Laboratory, Load Flexibility and Grid Integration. <https://datacenters.lbl.gov/load-flexibility-grid-integration>
5. Lawrence Berkeley National Laboratory and National Academy of Engineering, Integrating AI Data Centers with the Power Grid, 2026. <https://eta-publications.lbl.gov/publications/integrating-ai-data-centers-power>
6. Lawrence Berkeley National Laboratory, Large Load Literature Review, December 2025. https://eta-publications.lbl.gov/sites/default/files/2026-01/lbnl_illreview_dec_update_2025_final.pdf
7. US Department of Energy, Clean Energy Resources to Meet Data Center Electricity Demand. <https://www.energy.gov/oe/clean-energy-resources-meet-data-center-electricity-demand>
8. US Department of Energy, Rethinking Data Centers with AI-Driven Flexibility. <https://www.energy.gov/documents/pios202-25-11exhibits41to60>
9. US Department of Energy, Monitoring Oscillations from Large Data Centers, May 2026. <https://www.energy.gov/oe/articles/monitoring-oscillations-large-data-centers>
10. Electric Power Research Institute, DCFlex Demonstrations. <https://dcflex.epri.com/demonstrations>
11. Federal Energy Regulatory Commission, Technical Conference Regarding Large Loads Co-Located at Generating Facilities. <https://www.ferc.gov/news-events/events/commissioner-led-technical-conference-regarding-large-loads-co-located>
12. Federal Energy Regulatory Commission, Rulemaking RM26-4. <https://ferc.gov/rm26-4>
13. Federal Energy Regulatory Commission, FERC Directs PJM to Create New Rules for Large Loads, December 2025. <https://www.ferc.gov/news-events/news/fact-sheet-ferc-directs-nations-largest-grid-operator-create-new-rules-embrace>
14. Federal Energy Regulatory Commission, Targeted Action to Speed Large Load Integration, June 2026. <https://www.ferc.gov/news-events/news/ferc-launches-aggressive-targeted-action-speed-large-load-integration>
15. Federal Energy Regulatory Commission, Consumer Protection and Cost Recovery Agreements, September 2026. <https://www.ferc.gov/news-events/news/commissioner-rosners-concurrence-consumer-protection-and-cost-recovery-agreements>
16. Electric Reliability Council of Texas, Load Resource Participation. <https://www.ercot.com/services/programs/load/laar>
17. PJM Interconnection, Powering Reliability Through Market Design, May 2026. <https://www.pjm.com/-/media/DotCom/library/reports-notices/special-reports/2026/20260506-powering-reliability-through-market-design.pdf>
18. PJM Interconnection, Large Load Information. <https://www.pjm.com/markets-and-operations/large-load>
19. Virginia State Corporation Commission, Order in Dominion Energy Virginia Biennial Review and Large-Load Rate Case, 2025. <https://www.scc.virginia.gov/about-the-scc/newsreleases/release/scc-issues-order-on-dev-biennial-review-2025/scc-rules-in-dev-biennial-review-case.html>
20. Virginia State Corporation Commission, Large-Load Rate Facts. <https://www.scc.virginia.gov/about-the-scc/scc-facts/>
21. Georgia Public Service Commission, Data Center Factsheet. <https://psc.ga.gov/site/downloads/datacenterfactsheet.pdf>
22. North American Electric Reliability Corporation, 2025 Long-Term Reliability Assessment. https://www.nerc.com/pa/RAPA/ra/Reliability%20Assessments%20DL/NERC_Long%20Term%20Reliability%20Assessment_2025.pdf
23. National Institute of Standards and Technology, Cybersecurity Framework 2.0, 2024. <https://doi.org/10.6028/NIST.CSWP.29>
24. National Institute of Standards and Technology, Guide to Operational Technology Security, SP 800-82 Revision 3, 2023. <https://doi.org/10.6028/NIST.SP.800-82r3>
25. International Energy Agency, Energy and AI, 2025. <https://www.iea.org/reports/energy-and-ai>
26. US Department of Energy, Best Practices Guide for Energy-Efficient Data Center Design, July 2024. <https://www.energy.gov/sites/default/files/2024-07/best-practice-guide-data-center-design.pdf>
27. International Organization for Standardization, ISO/IEC 30134 Data Centre Key Performance Indicators. <https://www.iso.org/standard/63451.html>
28. International Valuation Standards Council, International Valuation Standards. <https://www.ivsc.org/standards/>
29. IFRS Foundation, IFRS 9 Financial Instruments. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-9-financial-instruments/>
30. IFRS Foundation, IFRS 15 Revenue from Contracts with Customers. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-15-revenue-from-contracts-with-customers/>

About the Author

Authored by Chennakeshav Adya

Independent Researcher

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.