

Move the Workload, Free the Grid: Revenue from Flexible AI Demand

A portfolio-finance framework for temporal and geographic compute flexibility

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Background. AI data-centre portfolios can contain deferrable workloads, spare compute, batteries, cooling flexibility, onsite generation and multiple locations. Electricity systems can value a verified change in demand during congestion or stress, while customers value completed computing service. Converting the same operating flexibility into dependable cash requires a framework that preserves both obligations. Objective. This paper develops a portfolio-finance method for measuring, contracting and monetising temporal and geographic AI-workload flexibility without double counting physical resources or compromising customer service. Method. It synthesises current official evidence from the International Energy Agency, the United States Department of Energy, the Federal Energy Regulatory Commission, the North American Electric Reliability Corporation and selected published programme frameworks. It then maps workload authority, electrical delivery, measurement, revenue and financing across a hypothetical multi-campus portfolio. Findings. The financeable product is a verified net change at a defined electrical boundary, delivered within a stated window and supported by workload rights, receiving-site capacity and customer-service controls. Gross megawatts moved between campuses can overstate grid value when receiving-site demand, battery discharge, rebound load and onsite generation are omitted. Revenue should be separated into availability, activation, congestion, capacity-release and customer-service components, with each component linked to an eligible programme or bilateral contract. Portfolio capital can recognise flexible-demand cash only after baseline, metering, dispatch, settlement, cyber, data, customer and non-performance terms are sufficiently firm. Implications. Operators, utilities, lenders and investors can use a flexibility ledger, matched-site test, revenue waterfall and downside framework to determine when workload mobility creates bankable value. Every numerical exhibit is hypothetical; it is not a market quotation, programme rule, forecast or description of a named project.

JEL Classification: G31, G32, L86, L94, O33, Q40, Q48

Keywords: AI data centres, flexible demand, workload shifting, congestion relief, demand response, grid services, compute orchestration, revenue stacking, data-centre finance, measurement and verification

This paper is an educational and structural analysis prepared for research purposes. It is not investment advice, an offer, or a solicitation. It contains no client information.

1. THE REMAINING COMMERCIAL QUESTION

Existing flexible-demand frameworks explain how a data centre can provide a measured grid service and how a bilateral agreement can contract that service. A further decision arises when an AI operator controls workloads across several campuses: what cash can be attributed to moving compute through time or geography, and what evidence makes that cash financeable?

The answer begins with two ledgers. The computing ledger records customer work, deadlines, service levels, data rights, hardware requirements and completion. The electricity ledger records demand at each connection point, dispatch, metering, baseline, rebound and settlement. Workload movement creates value only when both ledgers reconcile.

This paper develops that reconciliation. It focuses on portfolio revenue, capacity release and financing. It does not establish eligibility or payment under any named electricity programme. Those matters require current project-specific confirmation.

2. FLEXIBILITY IS A PORTFOLIO CAPABILITY

A single facility can delay work, adjust cooling, discharge storage or use permitted generation. A portfolio can also move eligible work to another time zone, electricity market or campus. This geographic option expands the decision set while creating receiving-site, network, data and customer dependencies.

The United States Department of Energy's data-centre flexibility taxonomy recognises timing and location as dimensions of flexibility and notes that control can be divided between facility owners and computing customers.[1] The International Energy Agency identifies operational flexibility and location as measures that can support data-centre integration into constrained power systems.[2][3]

Portfolio capability therefore belongs to an operating system, not to a server hall alone. It requires workload classification, spare compute, data movement, network capacity, compatible hardware, orchestration, cyber controls, energy signals and authority.

3. DEFINE THE DELIVERABLE PRODUCT

The grid product is a net demand change at a specified connection point relative to an approved reference. It has direction, magnitude, start, duration, notice, recovery and performance. The internal action can be temporal shifting, geographic transfer, throttling, battery discharge, cooling adjustment or generation.

The customer product is completed compute under agreed service, security and data conditions. A workload can be flexible for grid purposes only where the operator has authority to alter its time, location or power profile without breaching those conditions.

The portfolio contract should identify both products. Payment for grid flexibility cannot be supported solely by a scheduler forecast. Evidence must extend from dispatch signal through workload decision and facility response to meter and settlement.

4. DRAW THE DUAL-LEDGER ARCHITECTURE

The computing ledger begins with a job or service obligation. It records workload owner, deadline, permitted locations, hardware, checkpoint state, restart cost, data transfer and completion. The electricity ledger begins with a grid or price signal and records reference demand, instructed change, metered delivery, rebound and payment.

An orchestration layer links the ledgers. It selects eligible work, checks receiving capacity, applies customer and cyber constraints, executes the move and records the decision. A financial layer reconciles grid payment with incremental compute cost and customer consequences.




Figure 1. Dual-ledger architecture for financeable AI-demand flexibility

Generic framework; programme, customer and technical requirements must be verified for each portfolio.

5. CLASSIFY WORKLOADS BY OPERATING FREEDOM

The flexibility inventory should classify each workload by earliest start, deadline, duration, interruption tolerance, checkpoint capability, restart penalty, latency, residency, hardware, network and customer control. A broad description such as “training” or “batch” does not establish flexibility.

Some training stages may be deferrable but expensive to interrupt. Inference can be geographically distributed where latency, data and customer terms permit. Internal research can have wider windows than regulated or customer-facing applications. Backup, indexing and software builds can have their own constraints.

The inventory should be tested against actual scheduler data. Contracted flexibility should use the lower of operational evidence, customer authority, receiving capacity and programme qualification.

Table 1. Workload-flexibility evidence register

Workload class	Potential action	Evidence required	Principal constraint	Financeable quantity
Model training stage	delay, checkpoint or relocate	queue, checkpoint test, hardware match and deadline	lost progress and accelerator availability	net metered reduction within tested window
Batch inference	shift time or region	request profile, latency and receiving capacity	customer response and data location	verified portfolio demand change
Real-time inference	route eligible traffic	latency, resilience and customer permission	continuous service and network dependency	conservative tested tranche
Data processing	reschedule or split	completion deadline and transfer evidence	data movement and storage	change after transfer energy and rebound
Internal analytics	delay or relocate	owner approval and job records	business deadline	delivered meter response
Facility support	cooling, battery or generation action	engineering, state and permit evidence	resilience, degradation and operating limits	separately metered and non-duplicated response

Classification requires current scheduler, customer, data and technical evidence.

6. SEPARATE TEMPORAL FROM GEOGRAPHIC VALUE

Temporal shifting moves work within the same connection point. It can reduce demand during a defined interval and restore it later. Its economics depend on deadline, rebound, energy price, hardware utilisation and the value of completing work at a different time.

Geographic shifting reduces demand at one connection point and increases demand elsewhere. The system benefit depends on conditions in both locations. Moving 10 MW away from a constrained node can provide less than 10 MW of portfolio benefit if the receiving site is also constrained or if transfer and cooling increase total energy.

Contracts and models should keep the two products separate. Temporal delivery requires a recovery profile. Geographic delivery requires a matched-site record and rules for the receiving load.

7. CALCULATE NET PORTFOLIO RESPONSE

Gross reduction at the sending meter is only the first measurement. Net portfolio response should deduct incremental receiving-site demand, data-transfer energy and any simultaneous increase attributable to the move. Battery discharge and onsite generation should be recorded as separate resources where programmes distinguish them.

The portfolio can create local congestion relief even when global electricity use is unchanged or higher. Local value and portfolio energy cost therefore belong in separate model lines. An operator can earn a local service payment while incurring additional cost elsewhere.

The finance model should show sending-site settlement, receiving-site cost, transfer cost, recovery, customer impact and resource degradation. Cash remains financeable only after those items reconcile.

8. ESTABLISH WORKLOAD AUTHORITY

Authority can sit with the hyperscaler, tenant, colocation operator, software platform or end customer. Facility owners may control cooling and storage without controlling servers. Computing customers may control jobs without controlling the meter or market participation.

Each contracted tranche should identify the party authorised to schedule, move, interrupt and disclose operational data. Customer terms should address flexibility, location, data, service, liability and any benefit sharing. Silence should not be interpreted as permission.

The authority register should connect the electricity contract to underlying customer and operating rights. A portfolio cannot sell dependable flexibility beyond the weakest relevant authority.

9. PROTECT DATA RESIDENCY AND SOVEREIGNTY

Geographic movement can change the place where data is processed or stored. Applicable law, regulation, customer contracts and internal policy can restrict that movement. The assessment should cover input data, checkpoints, model weights, outputs, logs and backups.

A workload can be technically portable while remaining legally fixed. The orchestrator should enforce an approved location policy and retain evidence of each decision. Cross-border transfer, government access, encryption and subcontractor arrangements require qualified advice.

Revenue assumptions should exclude locations that are not authorised. A future sovereign-cloud arrangement can be modelled as an option until the required rights and infrastructure exist.

10. TREAT NETWORK CAPACITY AS A RESOURCE

Workload mobility depends on fibre, routing, latency, bandwidth, peering, data-transfer time and resilience. Large checkpoints or datasets can make a short electricity dispatch window operationally unusable. Network cost can exceed the grid payment for some workloads.

The receiving site also needs spare compute with compatible accelerators, software and storage. Spare electrical capacity without compatible compute cannot complete the work. Reserved compute can carry an opportunity cost when it displaces customer demand.

The availability test should therefore verify the entire route from scheduler to receiving compute and completed job. Electrical and digital capacity should be measured together.

11. DEFINE THE FLEXIBILITY PRODUCTS

A portfolio can offer day-ahead shifting, intraday reduction, fast curtailment, emergency interruption, geographic transfer, upward demand during surplus periods or a combination. Each product has different notice, response, duration, frequency and recovery.

Availability payment can compensate the reservation of qualified capacity. Activation payment can compensate delivered response. Congestion payment can reflect locational need under an eligible arrangement. A capacity-release payment can compensate a commitment that supports a connection or network-planning decision. Customer-service value can arise from lower energy cost or earlier access to constrained capacity.

The same capability should not be sold twice for overlapping intervals. The commitment register should reserve each resource by product, time and priority.

12. MAP ELIGIBLE REVENUE BEFORE FORECASTING IT

Technical capability does not create a payment route. Revenue requires an eligible programme, tariff, bilateral contract or internal cost allocation. Registration, metering, telemetry, aggregation, credit, dispatch and settlement conditions can apply.

FERC's market and large-load proceedings, DOE's flexibility work and NERC's computational-load programme demonstrate active institutional attention in the United States.[1][4][5] Their materials do not establish participation or price for a particular facility. Each regional and utility framework requires verification.

The revenue register should cite the governing document, counterparty, product, term, capacity, payment formula, performance, penalties and current status. An announced policy objective belongs outside contracted revenue.

13. BUILD THE REVENUE WATERFALL

Gross grid payment is reduced by aggregator or market fees, incremental electricity, receiving-site cost, data transfer, scheduler and control cost, battery degradation, generation fuel, customer compensation, recovery and tax. The remainder supports operating contribution and capital.

An internal energy saving should not be added to external settlement when both arise from the same avoided consumption without reconciliation. A customer payment for compute completion should remain in the computing case; only the incremental effect of flexibility belongs in the flexibility case.




Figure 2. Revenue waterfall from grid payment to financeable flexibility cash

Values are hypothetical index points and do not represent a tariff, market price or named portfolio.

14. DESIGN THE BASELINE

Settlement commonly depends on a reference for demand absent dispatch. The baseline can use historical intervals, schedules, forecasts or other programme methods. AI loads can grow quickly and vary with customer work, making baseline design material.

The method should address new facilities, ramping load, outages, maintenance, weather, model retraining, customer launches and prior flexibility. Adjustments require defined governance. An operator should not increase pre-event demand to inflate the apparent response.

Lenders need baseline transparency because settlement disputes affect cash. Source data, calculation code, versions, approvals and exceptions should be retained for each event.

15. MEASURE DELIVERY AND REBOUND

Delivery is the metered difference attributable to the approved action under the contract method. The meter interval, time synchronisation, losses and data quality should match the product. Receiving-site demand should be linked when geographic movement is claimed.

Rebound occurs when delayed work or storage recharge increases later demand. It can create cost or a second congestion event. The contract should define whether and how rebound affects performance or settlement.

The event report should reconcile instruction, action, sending demand, receiving demand, supporting resources, completed workload, rebound and payment. Missing telemetry should trigger an agreed estimation or dispute process.

16. PRESERVE CUSTOMER SERVICE

Customer-service limits should override economic dispatch according to an approved hierarchy. Critical inference, safety, cyber, regulated processing and contractual deadlines can remain protected. The operator should define who can suspend a response when service risk changes.

Service credits, lost revenue, customer churn and reputational damage can exceed a grid payment. The downside model should include those exposures and the limits of liability under customer and grid contracts.

Flexibility can also improve service by diversifying workload locations and exercising recovery pathways. That resilience value requires successful tests and should remain distinct from grid-service revenue.

17. AVOID DOUBLE COUNTING PHYSICAL RESOURCES

A battery can preserve IT load while reducing grid imports. Onsite generation can do the same. Workload shifting can reduce the load itself. A single metered response can combine all three, but the portfolio should attribute delivery by resource.

Priority rules should reserve energy and power for safety, ride-through, customer resilience and contracted grid service. A battery committed to two programmes in the same interval is not fully available to both. Generator use must comply with permits and customer commitments.

The resource ledger should track capacity, duration, state, maintenance, reservation and settlement. Financeable revenue uses the non-overlapping committed quantity.

18. CREATE THE MATCHED-SITE TEST

A geographic event should identify sending site, receiving site, workload tranche, start, duration and completion. The sending site must show a measurable reduction. The receiving site must show available compute, power and network capacity without breaching its own commitments.

The matched-site test should include both grids' conditions. A transfer can relieve one constrained region while adding demand to another. If the contract values only the sending node, the operator still needs to price the receiving consequence.

The test also captures failure modes: transfer delay, incompatible hardware, network congestion, data restriction, receiving outage and incomplete job. These risks determine the dependable tranche.

19. USE A PORTFOLIO CONTROL HIERARCHY

The hierarchy can begin with no-regret actions such as scheduling within approved windows, then use thermal or storage resources, then geographic transfer, and finally interruption where permitted. The order should reflect customer impact, cost, emissions, degradation and recovery.

Automated controls can execute within predefined limits. Human authority remains necessary for exceptions, new workload classes, customer overrides and material risk. Cyber controls should protect dispatch signals and prevent an external instruction from directly compromising computing systems.

Every action should be logged with source signal, policy version, decision, approver, command, meter result and customer outcome.

20. PRICE THE RECEIVING-SITE OPPORTUNITY COST

Receiving capacity may have an alternative use: serving a customer, running another job or preserving resilience. The cost of geographic flexibility includes the contribution forgone from that alternative. It also includes extra cooling, energy, network and hardware use.

Transfer pricing between portfolio entities should reflect the receiving service and applicable tax and regulatory requirements. An internal payment can allocate economics, but it does not create consolidated revenue.

The investment case should evaluate consolidated cash first, then allocate it. Entity-level financing follows enforceable contracts and cash movements.

21. RECOGNISE CAPACITY-RELEASE VALUE CAREFULLY

Flexible demand may help a utility accommodate a connection or defer an upgrade under a defined arrangement. The value can appear through a lower connection cost, earlier date, interruptible tariff, capacity payment or bilateral contract. It should not be assumed from technical flexibility alone.

Capacity-release claims require a network study or binding counterparty arrangement. The project should identify the period, conditions, dispatch rights and consequences of non-delivery.

Where flexibility accelerates campus revenue, that benefit belongs in the main campus counterfactual. A separate flexibility valuation should avoid adding the same acceleration twice.

22. MATCH CONTRACTS ACROSS THE PORTFOLIO

The sending-site agreement, receiving-site arrangement, customer contract, aggregator terms and internal orchestration rules should align. Conflicting dispatch instructions, service priorities or data obligations can make nominal capacity unavailable.

The portfolio agreement should allocate dispatch authority, measurement, settlement, cost, penalties, liability, force majeure, cyber incidents, data restrictions, maintenance and termination. Direct agreements can preserve essential rights for lenders.

The contract map should identify which entity receives grid payment and which entity incurs receiving cost. Cash transfer should follow documented service rather than an unsupported management allocation.

23. ALLOCATE NON-PERFORMANCE

Non-performance can arise from workload unavailability, scheduler error, customer override, receiving-site constraint, network failure, battery state, generator failure, meter error or external events. The contract should distinguish causes and remedies.

Penalties should be modelled with event frequency and downside concentration. Repeated failure can reduce future qualification or terminate participation. Customer remedies can arise simultaneously.

The operator should maintain a cure process and root-cause record. Financing covenants can require performance thresholds, reserves and remedial plans before distributions.

24. PROTECT AGAINST REBOUND CONCENTRATION

Several deferred workloads can become due at the same time. Uncontrolled recovery can create a new peak, increase tariffs or overload receiving resources. The scheduler should stagger recovery across sites and customer deadlines.

The revenue model should include recovery energy and demand charges. A product that pays for two hours of reduction can be uneconomic when the subsequent peak is expensive or customer completion requires overtime compute.

Recovery policy belongs in operating tests. It should also be visible to the grid counterparty where programme rules require it.

25. BUILD THE CAPITAL CASE

The required assets can include orchestration software, telemetry, meters, network, spare compute, storage controls, cyber systems and integration. Some are shared with core operations. Incremental

capital should include only the expenditure necessary for the flexibility product, with a documented allocation for shared assets.

Development equity funds studies, tests, customer changes and programme qualification. Equipment or software finance can support identifiable assets. Operating debt can recognise contracted net cash after performance history and counterparty review.

An assumed future grid payment should not fund committed capital. Stage-gated investment can stop when eligibility, measurement or customer authority fails.

26. USE THE HYPOTHETICAL PORTFOLIO CASE

Consider three hypothetical campuses with different grid conditions and spare compute. Site A faces a congestion event. Site B can receive some eligible work, while Site C has storage and renewable availability. The operator tests temporal delay, geographic transfer and physical support.

The worked case uses index values rather than currency or asserted megawatts. It demonstrates how gross sending-site reduction can exceed net portfolio response and how revenue can fall after receiving and recovery cost.

Every input is a hypothetical modelling assumption. It does not describe a programme, facility, customer or transaction.

27. TEST DISPATCH AND SETTLEMENT CASES

Table 2. Hypothetical portfolio dispatch and settlement cases

Case	Sending reduction index	Receiving increase	Physical support	Rebound	Net portfolio response	Revenue treatment
Temporal delay	100	0	0	70 later	100 during event	settlement less recovery cost
Geographic transfer	100	65	0	5	35 portfolio net	local congestion value plus receiving cost
Battery-supported	100	0	60	60 recharge	40 workload component	separate battery and workload attribution
Hybrid portfolio	100	35	25	25	40 portfolio net	resource-level settlement and no double count
Failed transfer	40	20	0	10	20 portfolio net	performance deduction and customer consequence

Index values are illustrative and do not represent programme rules, prices or operating results.

28. STRESS THE ECONOMICS

The principal variables include eligible workload share, event frequency, payment, receiving energy, data transfer, spare-compute opportunity cost, rebound, performance and customer impact. These variables can correlate. Congestion events can coincide with high receiving prices or limited portfolio headroom.

Figure 3 shows a hypothetical net-cash index across delivered performance and receiving-cost scenarios. It illustrates why gross payment alone cannot support an investment decision.




Figure 3. Hypothetical net-cash sensitivity to delivered performance and receiving-site cost
Index values are illustrative and do not represent market prices, forecasts or a named portfolio.

29. SIZE FINANCING TO DEPENDABLE NET CASH

Lenders should use contracted and tested net cash after receiving cost, event deductions, resource degradation and customer exposure. Event-driven revenue can be volatile and seasonal. Availability payments can be more stable when qualification and performance remain intact.

Debt sizing should test termination, rule change, lower event frequency, baseline dispute, customer override and site outage. A portfolio with several independent locations can diversify delivery, subject to correlated grid and network events.

Security can include receivables, accounts, contracts, software rights and equipment. The lender must verify whether programme registration and dispatch rights survive enforcement or require consent.

30. RECONCILE ENTITY CASH

Grid payment may accrue to a local facility while receiving cost arises in another entity or jurisdiction. Intercompany agreements should allocate service, data, energy and orchestration. Transfer pricing, tax, currency and withholding require qualified advice.

Consolidated value should be calculated before entity allocation. A local profit created by undercharging the receiving site is not portfolio value. Cash availability for debt depends on legal movement, covenants and minority interests.

The model should show each entity's receipts, costs, debt service, taxes, reserves and distributions, then reconcile to the group.

31. BUILD THE MEASUREMENT AND VERIFICATION OFFICE

The office maintains workload classification, customer rights, resource state, baselines, dispatch, meters, event reports, settlements and disputes. Operations, energy, finance, legal, cyber, data and customer teams require defined ownership.

Independent assurance can be used where contracts or financing require it. The assurance scope should cover the actual data chain and calculation rather than a general system description.

Version control is essential. A scheduler update, new customer, changed meter or revised baseline can alter the financeable product.

32. ESTABLISH CYBER AND OPERATIONAL CONTROLS

An electricity signal should enter the data-centre environment through authenticated, monitored and limited interfaces. It should not create unrestricted access to computing orchestration. Manual and automatic fail-safe states should be defined.

The operator should test false signals, delayed messages, conflicting instructions, loss of telemetry, receiving-site failure and recovery. Incident response should coordinate grid and customer obligations.

Cyber cost and downtime belong in the economics. Insurance and contractual liability require specialist review.

33. CONSIDER CARBON CLAIMS SEPARATELY

Moving work to a different time or place can change the electricity mix physically consumed. Contractual renewable attributes, location-based emissions, market-based reporting and customer claims are distinct. The evidence should identify the method and rights.

A lower-emissions schedule can have strategic value even without a grid payment. That value should not be represented as cash unless a customer or programme pays for it. Additional energy from transfer or recovery should remain visible.

Environmental claims require current data, consistent boundaries and qualified review. Flexibility revenue does not by itself establish an emissions reduction.

34. USE GEOGRAPHIC DIVERSIFICATION CAREFULLY

Multiple regions can diversify weather, grid and market exposure. They can also create correlated technology, network, customer and operator dependencies. A common orchestration failure can affect the entire portfolio.

The availability model should test simultaneous events, constrained receiving capacity and regional rules. It should not assume that every site can receive work during its own peak or emergency.

Portfolio diversification receives financing value only where rights, capacity and operating tests support it.

35. NEGOTIATE CUSTOMER BENEFIT SHARING

Customers can provide the workload authority that enables grid revenue. They may expect lower service prices, sustainability benefits, resilience or a share of payment. The agreement should state the benefit and the customer's obligations.

Benefit sharing should follow delivered economics after cost and penalties. A gross-revenue share can leave the operator with losses in high-cost events. Minimum service standards and opt-out conditions should be explicit.

Customer consent and reporting should avoid disclosing confidential grid or other-customer information.

36. PROTECT TRANSACTION VALUE

An acquisition of a flexible-compute portfolio can include projected grid revenue. The buyer should verify programme access, customer authority, measurement, receiving capacity, performance history and cost. A headline megawatt claim is insufficient.

Consideration can distinguish existing contracted cash from future qualification. Deferred price, earn-out or retention can link value to settled events and continuing eligibility. The purchase agreement should address baseline changes, customer loss and rule change.

The buyer should also identify whether orchestration software, data and key personnel transfer. Revenue dependent on a non-transferable affiliate capability requires a transition arrangement.

37. APPLY ACCOUNTING AND VALUATION DISCIPLINES

Contracts, software, customer relationships and financial instruments require analysis under the applicable accounting framework. IFRS 13 provides fair-value principles where fair value is required.[6] Recognition and classification require complete facts and qualified advice.

Income valuation should use supportable net cash and explicit renewal or rule-change assumptions. Market comparisons require similar programmes, rights, portfolio topology and maturity. Cost can inform replacement expenditure without establishing economic value.

The International Valuation Standards Council publishes standards for valuation practice.[7] Purpose, basis, date, scope, evidence and uncertainty should be documented.

38. USE STAGED APPROVAL GATES

Table 3. Staged flexibility-commercialisation gates

Gate	Decision	Minimum evidence	Capital released	Stop or reprice trigger
1	inventory flexibility	scheduler data, customer terms and site map	analysis and testing budget	insufficient workload authority
2	prove delivery	controlled events, meters and receiving-site tests	integration and telemetry	unreliable response or customer harm
3	qualify product	eligible programme or bilateral terms	registration, controls and contracting	no payment route or adverse liability
4	operate pilot	approved baseline, dispatch and settlement	limited operating capital	failed events, dispute or cyber weakness
5	scale portfolio	settled cash, diversified sites and governed commitments	expansion capital and working capital	double commitment or receiving constraint
6	finance cash flow	dependable contract, performance history and lender rights	operating debt or refinancing	termination, rule change or customer concentration

Governance template; all conditions require portfolio-specific evidence.

39. MAINTAIN THE FLEXIBILITY LEDGER

The ledger should show workload tranche, authority, sending site, receiving site, time window, electrical quantity, physical support, reservation, programme, price, event, meter result, rebound, customer outcome, settlement and dispute. Each field needs a source and owner.

The ledger prevents double commitment across programmes and resources. It also allows finance to reconcile gross payment to net cash and customer work.

Exceptions should trigger action. Missing receiving evidence, expired customer authority or failed performance should reduce the available quantity until resolved.

40. A 100-DAY EXECUTION PROGRAMME

Days one to twenty map workloads, sites, customer rights, meters, programmes and receiving capacity. The team establishes the dual-ledger data model.

Days twenty-one to forty run controlled temporal and geographic tests. Engineering, compute operations and customer teams validate delivery, completion and rebound.

Days forty-one to sixty obtain programme or bilateral terms, build baselines, design cyber controls and model net cash. Legal and finance map entity flows.

Days sixty-one to eighty operate pilot events, reconcile settlements and test failure cases. The team updates the dependable tranche.

Days eighty-one to one hundred approve scaled commitments, reserves, customer benefit sharing and financing. Actual timing depends on counterparties and portfolio maturity.

41. INVESTMENT-COMMITTEE CHECKLIST

Product. What exact demand change is sold at which connection point?

Workload. Which jobs can move, under whose authority and within what limits?

Portfolio. Which receiving site has compatible compute, power, network and data rights?

Measurement. How are baseline, meter, rebound and physical resources reconciled?

Revenue. Which current contract creates payment and what deductions apply?

Customer. Which service, data and liability provisions control dispatch?

Capital. Which incremental assets and working capital are required?

Downside. What happens under lower performance, higher receiving cost or programme change?

Finance. Which dependable net cash and rights support debt?

Governance. Who approves commitments, events, exceptions and reporting?

42. LIMITATIONS

Electricity programmes, grid conditions, data-centre technology, customer demand, data rules and financing conditions vary by jurisdiction and can change. Official system-level publications do not establish eligibility, price, revenue, authority or performance for a particular portfolio.

DOE, FERC and NERC materials apply within their stated United States purposes. IEA analysis is global and scenario-based. Accounting and valuation standards require application within their formal scope by qualified professionals.

Every quantity, index and case in the exhibits is hypothetical. No programme payment, customer right, meter result, financing offer, transaction value or realised outcome is claimed.

43. CONCLUSION

Flexible AI demand becomes financeable when computing and electricity evidence reconcile. The portfolio must complete customer work while producing a verified demand change at the contracted electrical boundary.

Temporal and geographic movement create different products. Geographic value requires a matched-site test and a net portfolio calculation. Batteries, onsite generation, cooling and workload actions need separate attribution to prevent double counting.

Revenue begins with an eligible contract and ends with net cash after receiving energy, transfer, recovery, degradation, customer and programme cost. A governed flexibility ledger, staged approvals and performance history can convert operational capability into an investable cash-flow proposition.

44. BOARD REPORTING SHOULD DISTINGUISH CAPABILITY, COMMITMENT AND CASH

The board dashboard should separate technically available flexibility, customer-authorised flexibility, programme-qualified flexibility, contractually committed flexibility, dispatched delivery, settled delivery and collected cash. These measures answer different questions. Combining them into one megawatt figure can conceal the point at which an operating capability fails to become revenue.

The dashboard should also show sending and receiving headroom, workload deadlines, physical-resource reservations, baseline exceptions, rebound, event performance, customer incidents, disputed settlements and net contribution. Each metric needs a definition, source, period and accountable owner. Forecast values should remain distinguishable from contract and settlement evidence.

Capital decisions should follow the evidence state. Software and integration expenditure can be approved for a controlled pilot before a long-term revenue contract exists, provided the board understands the stage and loss limit. Expansion capital should require repeated delivery and a credible route to net cash. Operating debt should require stronger evidence, including contractual term, counterparty, performance, termination and lender rights.

Collected cash is the final proof of commercial conversion. A successful dispatch without accepted settlement demonstrates technical performance while leaving revenue unverified. A settlement statement without collection leaves counterparty and timing exposure. The reporting chain should therefore connect each event from commitment to cash receipt and preserve unresolved differences.

This discipline also improves transaction readiness. A buyer or lender can reproduce the dependable tranche, identify the value assigned to uncontracted options and test whether portfolio claims survive customer, grid and receiving-site diligence.

FREQUENTLY ASKED QUESTIONS

How is this framework different from a demand-response agreement?

It focuses on portfolio economics and financing after a response product exists. It reconciles workload movement, receiving-site cost, capacity release and net cash across several campuses.

Does moving 10 MW of compute create 10 MW of grid value?

Local sending-site relief can be 10 MW while net portfolio response is lower after receiving demand, transfer energy, physical support and rebound. Contract measurement and consolidated economics should both be shown.

Can geographic shifting support debt?

It can contribute where payment is contracted, workload rights and receiving capacity are verified, performance is demonstrated and net cash remains dependable under downside cases.

How should batteries and onsite generation be treated?

They should be recorded as separate physical resources with their own capacity, state, cost and restrictions. The same metered reduction should not be attributed fully to workload movement and a physical resource.

What is capacity-release value?

It is an economic benefit created where a verified flexibility arrangement supports a connection, tariff, network decision or earlier service. It requires counterparty evidence and must not duplicate the main campus acceleration value.

What is the core control?

The flexibility ledger connects customer authority, workload action, sending and receiving sites, physical resources, meter delivery, rebound, settlement and net cash for every committed tranche.

REFERENCES

United States Department of Energy. "Powering AI and Data Center Infrastructure: Recommendations," including Appendix A, Taxonomy for Data Center Power Flexibility. July 2024. <https://www.energy.gov/sites/default/files/2024-11/Powering%20AI%20and%20Data%20Center%20Infrastructure%20Recommendations%20July%202024.pdf>

International Energy Agency. Energy and AI, including "AI and energy security," "Energy demand from AI" and the executive summary. Paris, 2025. <https://www.iea.org/reports/energy-and-ai>

International Energy Agency. "Scaling Up Demand Flexibility," executive summary. Paris, 2026. <https://www.iea.org/reports/scaling-up-demand-flexibility/executive-summary>

Federal Energy Regulatory Commission. State of the Markets 2025, including large-load and data-centre flexibility discussion. March 2026. https://www.ferc.gov/sites/default/files/2026-03/26_State-of-the-Market_0324_1430.pdf

North American Electric Reliability Corporation. "Large Loads Action Plan," current computational-load resources. 2026. <https://www.nerc.com/initiatives/large-loads-action-plan>

IFRS Foundation. "IFRS 13 Fair Value Measurement." <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/>

International Valuation Standards Council. "International Valuation Standards." <https://www.ivsc.org/standards/>

International Energy Agency. Electricity 2026, "Flexibility." Paris, 2026. <https://www.iea.org/reports/electricity-2026/flexibility>

United States Department of Energy, Office of Electricity. "Microgrids, Large Electric Loads & Grid Support: How to Leverage Microgrids to Support Utilities and Large Load Customers." 2026. <https://www.energy.gov/oe/articles/microgrids-large-electric-loads-grid-support-how-leverage-microgrids-support-utilities>

International Energy Agency. "The Value of Demand Flexibility," executive summary. Paris, 2025. <https://www.iea.org/reports/the-value-of-demand-flexibility/executive-summary>

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.

Matchpoint profile: <https://www.matchpoint-partners.com/team/ck-adya.html> LinkedIn: <https://www.linkedin.com/in/ckadya/>

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.