The Compute Tariff: Pricing Power, Cooling and Network into an AI Capacity Contract

A Contract and Financing Framework for Bankable AI Infrastructure Revenue

Authored by Chennakeshav Adya
Independent Researcher
August 2026

Abstract

Artificial-intelligence infrastructure is frequently described through megawatts, rack density, accelerators and headline availability. A customer buys a delivered computing service. The commercial contract must connect reserved physical capacity to usable compute, cooling, network, security, measurement, service quality, change and payment. That translation determines whether demand can support development capital, equipment finance, working capital and long-duration infrastructure investment.

This paper develops a board framework for designing and underwriting an AI capacity contract. It defines the service boundary, separates reservation, consumption and outcome charges, maps power and thermal dependencies, creates a network and latency schedule, specifies acceptance and ramp, governs accelerators and technology refresh, and links service-level failure to measured remedies. It also covers metering, billing, customer credit, collateral, curtailment, force majeure, data portability, termination, step-in, lender diligence and portfolio concentration.

The evidence base includes the International Energy Agency's 2025 Energy and AI report and April 2026 update, the United States Department of Energy and Lawrence Berkeley National Laboratory data-centre energy work, Commission Delegated Regulation (EU) 2024/1364, National Institute of Standards and Technology cloud-service metric guidance, and Open Compute Project guidance for liquid distribution and advanced cooling. These sources establish relevant demand, measurement and engineering considerations. They do not prescribe a universal commercial tariff or establish bankability for any project.

Six original figures and six implementation tables support an illustrative capacity contract and a 120-day commercialisation office. Every capacity, price, utilisation, efficiency, availability, power, water, cost, penalty, credit and valuation input is a management assumption created solely to demonstrate the method. It is not a customer offer, market quote, forecast, engineering guarantee or financing commitment. Live transactions require current engineering, technology, cyber-security, energy, environmental, legal, regulatory, tax, accounting, insurance, credit and financing advice in each relevant jurisdiction.

JEL Classification: G32, L86, L94, O33, Q41

Keywords: AI capacity contract, compute tariff, data centre, power, liquid cooling, network, service levels, project finance

1. Price the delivered service boundary

An AI capacity contract should begin with the customer outcome and work backwards through the technical stack. The buyer may require training capacity, inference throughput, dedicated accelerator clusters, a private cloud environment, sovereign processing, disaster recovery or an agreed combination. Each product needs a defined service boundary.

The boundary identifies the compute equipment, orchestration layer, storage, network, electrical path, thermal path, facility, security controls and operating services included in the charge. It also lists customer-provided elements. A contract that promises accelerator access while leaving network, storage or cooling performance undefined can create a service that exists technically and fails commercially.

The contract unit should connect four layers. Reserved capacity defines what the provider holds for the customer. Activated capacity defines what has passed acceptance and can carry workload. Consumed capacity measures actual use. Delivered service measures the quality and performance of that use. Revenue recognition, billing and remedies should follow the applicable accounting policy and executed contract.

The International Energy Agency estimates that data centres consumed about 415 terawatt-hours in 2024 and projects about 945 terawatt-hours in 2030 in its 2025 base case. Its April 2026 update reports 17 per cent growth in data-centre electricity demand in 2025 and describes constraints in grids, approvals, turbines, transformers, chips and IT equipment. Those system-level findings support rigorous capacity definition. They do not establish demand for a specific facility.

Figure 1. Delivered AI capacity service boundary

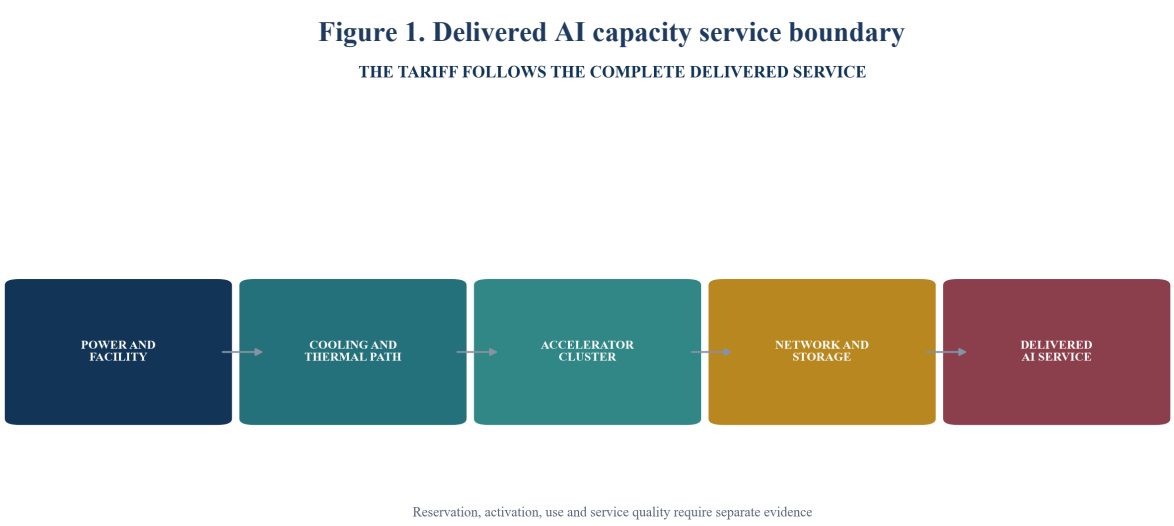

The contracted unit connects physical inputs, digital infrastructure, controlled operation and the customer service.

2. Define the capacity product before price

A tariff cannot repair an ambiguous product. The provider should issue a capacity specification that a customer, engineer, operator, lender and auditor can interpret consistently. The document defines accelerator type and approved substitutes, quantity, interconnect, memory, storage, network ports, rack power, cooling class, software layer, physical location, security perimeter, support and service period.

Capacity can be sold as a dedicated cluster, reserved accelerator quantity, minimum compute hours, power-backed rack allocation, workload throughput or managed outcome. These units create different risks. A dedicated cluster creates asset and obsolescence exposure. A compute-hour product creates utilisation and scheduling exposure. An outcome product adds application quality, model and data dependencies.

The specification should separate firm, interruptible and burst capacity. Firm capacity receives a reservation commitment, defined maintenance treatment and agreed service level. Interruptible capacity can be curtailed under specified triggers and may receive a lower price. Burst capacity depends on availability and a price or index at the time of use.

Substitution rights require precision. A nominally newer accelerator may have different memory, interconnect, software support, power density or performance on the customer workload. Equivalent capacity should be tested against an agreed benchmark and quality envelope. Customer consent, notice and price adjustment can apply when substitution changes the economic service.

Table 1. AI capacity product schedule

Field	Contract definition	Evidence at acceptance	Continuing measure
reserved capacity	named cluster, quantity and reservation window	inventory and allocation record	reservation ledger
accelerator service	model, memory, interconnect and approved substitute	configuration and benchmark	telemetry and change log
electrical service	firm, interruptible and burst power boundary	energisation and load-bank test	meter and event log
thermal service	cooling class, supply conditions and heat-removal duty	integrated systems test	temperature, flow and alarms
network service	ports, bandwidth, routes, latency boundary and diversity	end-to-end test	network telemetry
storage service	type, usable capacity, throughput, durability and location	performance and access test	utilisation and incident record
operations	support, maintenance, security and escalation	readiness certificate	ticket, control and maintenance data

Every field should be supported by an accepted design, test or operating record.

3. Use a three-part tariff

A practical tariff can separate reservation, consumption and service components. The reservation charge pays for capacity held available to the customer. It supports fixed facility, equipment, financing and operating commitments. The consumption charge follows measured use and passes through variable energy, network, software or service cost where appropriate. The service component rewards managed performance or an agreed outcome.

The allocation should follow controllable risk. A provider financing dedicated equipment needs enough committed revenue to cover debt service, operating cost, maintenance, refresh reserve and

an approved return under downside utilisation. A customer needs protection against paying for capacity that never becomes ready or cannot meet the agreed workload.

Minimum commitments can be phased. An initial reservation covers a defined cluster. Additional tranches activate after conditions such as facility readiness, customer demand approval, equipment delivery, acceptance tests or financing. Take-or-pay language requires legal and accounting review. Contracted revenue is economically useful only when the obligation, conditions, exclusions, credit and termination rights are understood.

Indexation should follow cost exposure. Energy can reference an agreed tariff, meter or price formula. Network, software or licence costs can follow verified third-party changes. A general inflation index can apply to appropriate operating components. Accelerator economics should reflect the refresh and residual-value plan rather than an automatic inflation escalator.

Figure 2. Three-part compute tariff

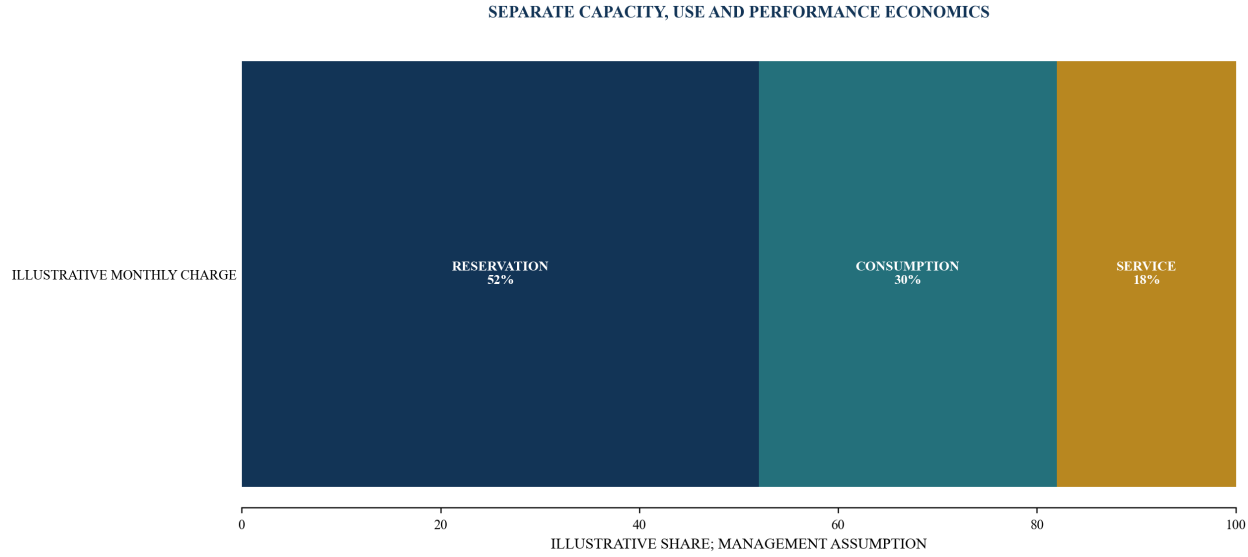

Reservation, measured use and delivered service address different economic risks.

4. Convert power into a contract schedule

Power is an input to the service and a source of delivery risk. The contract schedule should define the electrical boundary, maximum demand, ramp, energy measurement, quality, redundancy, outage treatment, backup limits, curtailment, maintenance and responsibility for utility or on-site supply.

Nameplate power is insufficient. The provider should reconcile contracted utility capacity, available facility capacity, protected reserve, conversion losses, cooling load, network and storage load, customer IT load, simultaneous use and expansion. The schedule should distinguish capacity that is contracted, connected, energised, commissioned, accepted and available for customer use.

The United States Department of Energy describes data-centre demand as rapidly growing, regionally concentrated and often requiring firm power. Its December 2024 summary of Lawrence Berkeley National Laboratory work estimated that US data centres could consume 6.7 to 12 per cent of national electricity by 2028. The current DOE resource hub cites a later 2030 central estimate of 11.8 per cent with a 9.5 to 15.3 per cent range. These are system scenarios. Project underwriting requires utility, site and equipment evidence.

Curtailment can create value when the workload, customer agreement and operating design permit flexibility. The contract must define trigger authority, maximum events, notice, duration, restoration, protected workloads, measurement and compensation. A provider should avoid selling the same firm capacity to a customer and a grid programme without a reconciled priority waterfall.

5. Contract the thermal envelope

High-density AI equipment turns power into heat at a concentrated point. The cooling schedule should state the technology, facility water and technology cooling boundaries, supply and return conditions, flow, pressure, water quality, heat-removal capacity, redundancy, leak detection, thermal ride-through, maintenance and alarms.

The Open Compute Project's Advanced Cooling Facilities guidance addresses facility water systems, liquid distribution loops, piping, routing, connections, temperature control, isolation, thermal ride-through, commissioning, procedures and service-level considerations. Its guidance supports disciplined interface design. The project engineer and equipment requirements govern the live design.

The customer and provider should agree what happens when equipment specifications change. A new accelerator generation can require different rack power, coolant temperature, flow or connection. The change process should test facility capability, cost, schedule, interoperability, warranty, safety and downtime before the provider promises compatibility.

Cooling performance should connect to billing and service. A facility efficiency metric such as power usage effectiveness can support operating analysis. It does not directly measure customer workload output. Water usage effectiveness can expose water intensity. Each metric needs a defined boundary, period, data source and treatment of exceptional conditions.

Table 2. Power and cooling interface schedule

Interface	Definition	Provider evidence	Customer dependency	Failure treatment
utility capacity	contracted supply and connection point	executed agreement and meter	approved ramp	delay or relief under defined cause
facility power	usable capacity after infrastructure load	capacity reconciliation	rack forecast	reservation adjustment
power quality	voltage, frequency and event boundary	monitoring and event data	compliant equipment	incident classification
cooling duty	supported heat load by zone and rack	design and integrated test	declared equipment load	capacity restriction or remedy
liquid interface	temperature, flow, pressure and chemistry	commissioning data	compatible technology loop	controlled shutdown and cure
thermal ride-through	supported duration after defined event	witnessed test	workload response	continuity plan
water boundary	source, quality, use and restriction	permits, meters and plan	agreed operating envelope	conservation or curtailment protocol

Thresholds and remedies depend on the engineered system and executed contract.

6. Define network as part of compute capacity

AI capacity without usable connectivity can become stranded equipment. The network schedule should define the service demarcation, ports, bandwidth, latency measurement points, jitter, packet loss, route diversity, internet and private connectivity, cloud on-ramps, cross-connects, data-transfer charges, maintenance, monitoring and incident escalation.

Latency requires an end-to-end boundary. A facility can control its internal network and selected carrier interfaces. The customer can control application design, data location and external routes. The contract should allocate each segment and define which measurement supports a service claim.

Bandwidth commitments require sustained and burst treatment. The provider may reserve physical ports while upstream capacity is shared. Oversubscription, traffic shaping and congestion management should be transparent within the agreed product. Egress charges can materially change workload economics and should be visible in the tariff model.

Route diversity should be evidenced physically. Two logical services can share a duct, exchange, building entrance or upstream dependency. The diligence file should include route maps, carrier commitments, demarcation records, failure domains, maintenance coordination and restoration priorities.

Data location and sovereignty can affect network architecture. The contract should identify the facility, approved replication locations, remote support, telemetry flows and customer-controlled encryption. Current legal and cyber-security advice should confirm the relevant requirements.

7. Use reproducible service metrics

A service-level commitment is only as strong as its metric. The National Institute of Standards and Technology's cloud-service metric framework describes the definition, unit, rules, measurement constraints and values required to understand a service property. NIST connects metrics to selection, service agreements, monitoring, accounting and audit.

The schedule should define the numerator, denominator, observation period, measurement point, clock, exclusions, source system, aggregation, evidence retention, dispute process and remedy for each metric. Availability can mean facility power, cluster accessibility, scheduler acceptance or successful workload completion. Each produces a different result.

Planned maintenance should have a notice period, annual or quarterly allowance, permitted window and overrun treatment. Emergency maintenance should have a separate definition and governance path. Broad exclusions can remove the economic meaning of an availability promise.

Remedies should follow the customer's loss mechanism and the provider's controllable failure. Service credits can apply automatically above an evidence threshold. Repeated or severe failure can trigger remediation plans, enhanced reporting, termination or transition support. Legal advice should address limitation of liability, sole-remedy language and enforceability.

Figure 3. Service-level metric chain

A SERVICE CREDIT BEGINS WITH A REPRODUCIBLE MEASURE

Boundary, period, source, exclusions and dispute process make the result auditable

A remedy should follow a reproducible metric and a controlled evidence record.

8. Make acceptance a revenue gate

Delivery should pass objective acceptance before committed billing begins. The acceptance plan covers facility, power, cooling, network, hardware, storage, software, security, monitoring, operations and the customer workload or agreed representative test.

Testing should proceed from component to integrated system. Equipment inventory and configuration come first. Electrical and thermal systems are tested through expected load and defined fault scenarios. Network tests confirm routes and performance. Cluster tests confirm health, orchestration, storage and workload execution. Security tests confirm access, segmentation, logging and approved controls.

The contract should distinguish material defects, minor punch-list items and customer-caused delay. A material defect prevents acceptance. A minor item can remain under a timed cure plan when it does not impair the agreed service. Customer dependencies such as data, code, access, approvals or test workloads need due dates and consequences.

Deemed acceptance can create disputes when the customer cannot run the test or the criteria are incomplete. A controlled mechanism can use independent evidence, a defined response period and an escalation process. Revenue accounting should follow the executed arrangement and applicable standards.

The acceptance certificate should identify the exact capacity tranche, configuration, service start date, outstanding items and approvers. That record becomes the basis for billing, warranties, maintenance and financing draw conditions.

Table 3. Capacity acceptance matrix

Workstream	Test	Evidence	Acceptance authority	Open-item rule
electrical	load, transfer and defined fault events	witnessed results and alarms	commissioning authority	material failure blocks tranche
thermal	full-load heat rejection and ride-through	temperatures, flow and event data	engineer and operator	bounded minor cure list
hardware	inventory, health and configuration	serial and telemetry record	provider and customer	failed unit replaced
network	bandwidth, latency, loss and route tests	end-to-end test file	network owners	route defect classified
platform	scheduler, storage, identity and logging	functional results	service owner	severity-based cure
workload	agreed representative benchmark	output, quality and performance	named customer approver	variance process
operations	support, monitoring and incident drill	tickets and drill record	operating committee	readiness condition

The exact tests and authorities should follow the customer product and engineered design.

9. Govern the ramp and capacity option

AI demand can grow quickly and unevenly. A phased contract should define initial capacity, committed ramp, customer options, provider expansion rights and the conditions for each tranche. Capacity reserved too early can depress utilisation. Capacity promised without secured power, cooling or equipment can create a delivery exposure.

An option has an economic cost. The provider holds land, power, rooms, equipment slots, network or procurement capacity for a customer. The option charge should reflect the reserved resource, expiry, exclusivity and probability that another buyer could use it. A refundable deposit, option fee, development contribution or reservation charge can allocate this exposure.

The capacity plan should identify long-lead dependencies for every tranche. Transformers, switchgear, generators, cooling equipment, accelerators and network components can have different procurement paths. The IEA's April 2026 update identifies tightening supply chains among the constraints affecting expansion. The live procurement file remains the project evidence.

Customer forecast obligations improve coordination. A rolling forecast can distinguish binding, expected and indicative demand. Variance bands, notice and priority rules guide scheduling. The provider should reconcile aggregate customer options to physical deliverability and financing capacity.

Expansion can require permit, utility and lender approval. The contract should avoid implying that a commercial option overrides third-party authority. Conditions precedent and long-stop dates should make this boundary clear.

10. Build technology refresh into the original economics

Accelerator economics can change during a long customer term. A contract should define the initial technology, useful service period, maintenance, approved substitutions, refresh triggers, benchmark method, migration, residual equipment, data handling, downtime and price reset.

Refresh can be mandatory, optional or performance-driven. A mandatory cycle gives predictability and may force capital before the customer values it. An optional cycle preserves choice and can create negotiation at the point of dependency. A performance-driven cycle uses an agreed workload and threshold, requiring stable tests and governance.

The provider should maintain a refresh reserve or financing plan. The economic model includes equipment purchase, installation, testing, lost service time, migration support, residual proceeds, decommissioning and disposal. Customer contributions and renewed term can support the investment.

Benchmark drift is a material risk. Software versions, precision, model architecture, data, batching and optimisation can change performance. The contract should preserve the benchmark package, test environment, quality target and review process. A raw peak-performance specification rarely represents the delivered workload.

Technology change also affects power, cooling and network. The refresh decision should re-open the full interface schedule. Compatibility should be demonstrated before commitment.

Figure 4. Capacity and technology lifecycle

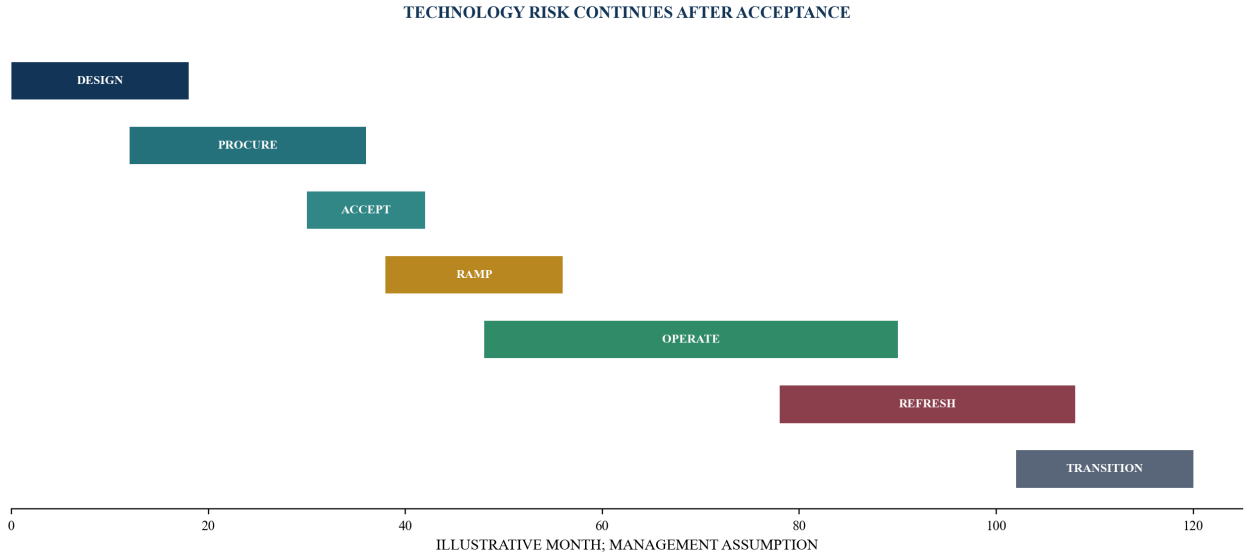

Contract, operations and financing should address every transition before the original capacity is committed.

11. Reconcile metering to invoice and cash

The revenue system should connect contract, asset, telemetry, tariff, invoice, receivable and cash. Each charge needs a source, calculation owner, review, customer evidence and correction process. Manual spreadsheets can support an initial controlled process. Scale requires stable identifiers and reconciliations.

The reservation ledger records customer, capacity tranche, location, configuration, start, end, status and price. The usage ledger records the agreed consumption meter and period. The service

ledger records performance, exclusions, incidents and credits. The invoice engine applies the approved tariff and tax treatment.

Meters need governance. The contract identifies authoritative sources, sampling interval, time zone, missing-data method, reset, calibration where relevant, retention and customer access. A provider-side dashboard can improve transparency. It should not replace the underlying audit record.

Disputes should be bounded. The customer raises a specified issue within an agreed period and pays undisputed amounts. The provider investigates against retained evidence. Corrections enter a controlled billing period. Repeated measurement defects can trigger a root-cause plan.

Cash conversion belongs in the board dashboard. Contracted revenue can coexist with delayed acceptance, billing disputes, credits and weak collection. The operating model should report reserved value, accepted recurring revenue, billed revenue, credits, receivables, cash and concentration separately.

Table 4. Meter-to-cash control ledger

Ledger	Key record	Control	Reconciliation	Decision use
reservation	customer and capacity tranche	approved capacity allocation	physical and contracted capacity	forward commitment
activation	accepted service and start date	signed acceptance	asset and billing start	recurring revenue gate
consumption	meter, unit and period	source and completeness check	telemetry to invoice quantity	variable revenue
service	performance and exclusion data	incident approval	metric to credit	quality economics
invoice	tariff calculation and tax	maker-checker review	contract to bill	receivable creation
collection	receipt and allocation	bank reconciliation	invoice to cash	realised conversion
correction	dispute and approved adjustment	authority and audit trail	correction to source	defect management

Definitions and controls should follow the executed contract and applicable accounting policy.

12. Underwrite the customer and security package

A long-term capacity contract transfers credit risk to the provider and its financiers. Diligence should cover customer identity, ownership, financial condition, funding, business model, demand evidence, authority, sanctions and financial-crime requirements. A startup, large technology company, government entity and project vehicle require different analysis.

The credit package can include deposits, advance payment, parent support, letter of credit, guarantee, pre-funded reserve, termination payment, equipment title rights or staged capacity. Legal advice should confirm enforceability, perfection, insolvency treatment and jurisdiction.

The security package should correspond to exposure. A dedicated accelerator purchase creates a larger early risk than a short shared-capacity agreement. Credit support can step down after acceptance, payment history or contracted milestones. It can step up after downgrade, late payment or adverse change when the contract permits.

Customer concentration requires a portfolio view. One anchor buyer can make initial financing possible and create renewal, termination and pricing exposure. The board should model customer failure, delayed ramp, partial renewal and replacement time. Marketing claims about contracted capacity should reconcile to termination rights and credit quality.

The provider's financiers may require assignment, notice, cure, direct agreements or step-in rights. Customer consent should be addressed during contract negotiation. These rights should preserve service continuity and the customer's data and security protections.

13. Design remedies around controllable failure

Remedies allocate economic consequences after failure. The contract should classify service failure, delayed delivery, capacity shortfall, chronic underperformance, security incident, data breach, payment default, customer-caused delay and force majeure. Each category needs evidence, notice, cure, mitigation and escalation.

Service credits can compensate for measured degradation. A credit schedule can rise by severity and frequency. A cap prevents unlimited routine credits and should not obscure separate rights for serious events. The calculation period, affected charge and carry-forward treatment require definition.

Delay remedies should reflect the capacity tranche and critical path. Liquidated damages require jurisdiction-specific advice and a defensible relationship to loss. Alternative treatment can include delayed billing, additional capacity, option extension or termination after a long-stop date.

Chronic failure matters more than one isolated incident. A rolling threshold can trigger executive review, independent technical assessment, a remediation plan, additional reporting and eventual termination. The provider should retain the right to cure and manage system integrity.

Customer misuse, unapproved configurations or workload behaviour can damage performance. Acceptable-use rules should be precise and proportionate. Suspension should follow defined severity, notice and emergency rights. The record should preserve why action was taken.

Figure 5. Failure-to-remedy ladder

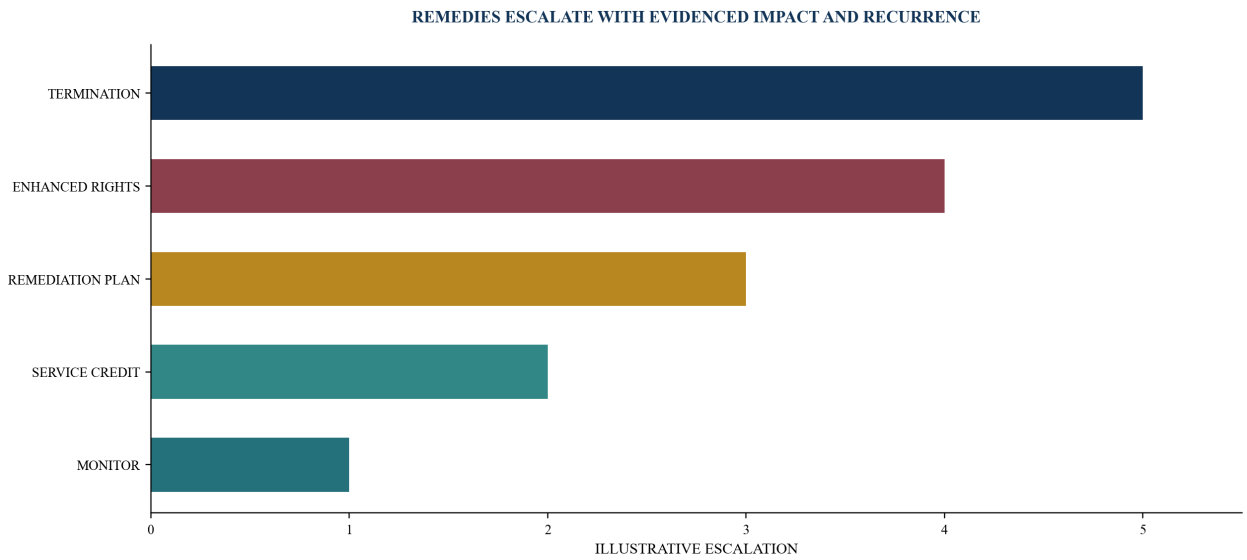

Severity, recurrence and control determine the contractual response.

14. Separate sustainability measures from service claims

Energy, water, renewable supply, refrigerants and waste heat can affect cost, permits, customer procurement and disclosure. The contract should identify which sustainability measures are supplied, how they are calculated, which party provides data and whether they influence price or service.

Commission Delegated Regulation (EU) 2024/1364 establishes information, measurement methods and sustainability indicators for specified reporting data centres with installed IT power demand of at least 500 kilowatts. It includes energy consumption, power usage effectiveness, water usage effectiveness, energy reuse factor and renewable energy factor. The regulation applies within its scope and uses defined reporting boundaries.

A tariff can pass through verified energy cost, reward load flexibility or price a defined renewable product. Contract language should avoid presenting certificates, power purchase agreements, on-site generation and hourly matching as interchangeable. The claim should follow the instrument and evidence.

Water cost and availability can affect the cooling choice. The US Department of Energy defines water usage effectiveness as annual site water use divided by IT-equipment energy use. A lower water value can coincide with different energy, capital or thermal consequences. Design decisions require site-specific engineering.

Customers can request emissions and resource information. The provider should disclose boundary, period, method, source, coverage and limitation. An operational metric should remain separate from a customer product guarantee unless the contract makes the connection explicit.

15. Create a transparent illustrative tariff model

A board model should use one coherent capacity case and disclose every assumption. Consider an illustrative 10-megawatt IT service supporting high-density accelerator capacity. The facility is assumed to ramp from 35 per cent accepted utilisation in year one to 82 per cent in year four. The numbers demonstrate the structure and do not describe an observed facility.

Assume a monthly reservation charge of USD 190 per available kilowatt, an energy pass-through based on metered use and a USD 22 per kilowatt managed-service component at full service. Assume power usage effectiveness of 1.28 initially, improving to 1.22, and an illustrative blended energy price of USD 0.085 per kilowatt-hour. All are management assumptions.

The model calculates contracted reservation revenue, variable energy recovery, service revenue, credits, direct facility cost, equipment cost, network, software, operations, maintenance and refresh reserve. Financing cash flow remains separate from operating contribution.

Downside cases reduce customer ramp, increase energy cost, delay acceptance, increase PUE, apply service credits and reduce renewal. The board should examine covenant and liquidity consequences. A model that survives only through uncommitted customer options does not establish bankability.

No approved observed Matchpoint or client revenue, customer pricing, capacity, utilisation, cost saving, mandate fee or collected revenue was supplied for this framework. Attributed monetary value remains USD 0 until supported by signed and paid mandate evidence or approved client

records.

Table 5. Illustrative compute-tariff assumptions

Assumption	Base case	Downside	Evidence required for live model
IT service capacity	10 MW	10 MW	accepted design and capacity reconciliation
year-one accepted utilisation	35%	22%	executed customer ramp and acceptance plan
monthly reservation charge	USD 190/kW	USD 170/kW	executed customer tariff
managed-service charge	USD 22/kW	USD 16/kW	service scope and contract
energy price	USD 0.085/kWh	USD 0.115/kWh	tariff, hedge and meter boundary
initial PUE	1.28	1.38	engineered model and operating test
annual service credits	1% of charges	6% of charges	service performance and credit schedule
renewal at initial term	75%	40%	customer decision and replacement pipeline

All values are management assumptions for method demonstration and require replacement with verified project evidence.

16. Translate the contract into financing evidence

Financiers need evidence that revenue, asset, cost and remedy structures align. The diligence package should include executed contracts, customer credit, capacity schedules, acceptance criteria, tariff formulas, termination rights, credit support, technology plan, equipment title, power and cooling evidence, network commitments, insurance and operating capability.

Revenue quality depends on enforceable obligation and deliverability. A long term can support financing while increasing technology and performance exposure. A take-or-pay provision can strengthen fixed revenue while remaining subject to conditions, exclusions, customer remedies and credit. Diligence should summarise the full arrangement.

The base case should match contracted capacity and current evidence. Customer options belong in a separate case until exercised. Unaccepted capacity should not be treated as operating revenue. Energy pass-through should be separated from margin. Service credits and termination should be modelled through the contract definitions.

Lenders can require reserve accounts, maintenance covenants, technology-refresh tests, concentration limits, additional customer commitments, completion support and cash sweeps. The structure depends on project, sponsor, jurisdiction and financing product.

The contract should support transition after enforcement. Direct agreements can give financiers notice and cure rights while preserving the customer's service, confidentiality, cyber-security and data protections. Current legal advice is essential.

Figure 6. Contract-to-financing evidence bridge

THE CAPACITY CONTRACT CONNECTS ASSET AND CASH

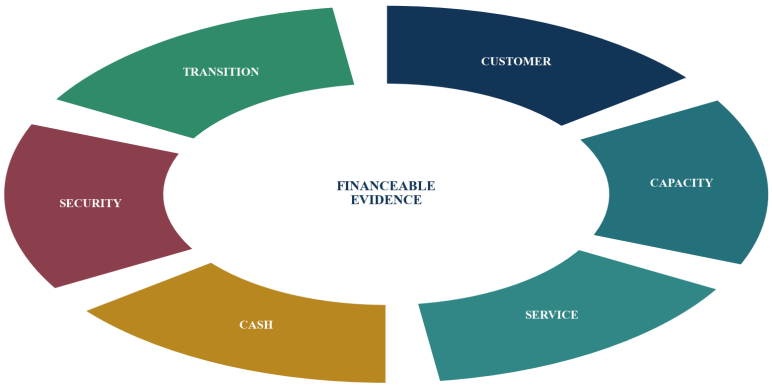

Bankability depends on alignment among demand, delivery, cash and downside rights.

17. Govern concentration and portfolio capacity

A provider should manage its whole portfolio rather than negotiate each customer in isolation. The portfolio register maps firm reservations, interruptible rights, options, physical capacity, power, cooling, network, equipment, credit support, sector, geography and renewal date.

Capacity overlap is a core control. The same accelerator, rack, power block or network headroom should not support multiple firm obligations. Options, burst rights, maintenance reserve and recovery capacity also consume deliverability. A weekly reconciliation can identify shortfall before it becomes a contract breach.

Customer concentration should be measured by accepted recurring revenue, committed reservation, receivable and replacement difficulty. Two customers with equal revenue can create different risk when one uses standard shared capacity and the other requires dedicated technology, custom networking and sovereign controls.

Renewal clustering can create a refinancing problem. The board should track term, notice, repricing, refresh and termination dates. Staggered maturity, extension options and early renewal discussions can improve visibility. Each extension should be tested against equipment and facility life.

Portfolio pricing should also reflect scarcity. A constrained power or high-density cooling zone can have different value from general capacity. The allocation committee should compare contribution, credit, strategic fit, contract strength and optionality using an approved scorecard.

18. Build security, data and continuity into service design

The service architecture should allocate physical security, identity, privileged access, network segmentation, vulnerability management, logging, incident response, data protection, encryption, key control, backup, recovery, deletion and audit support. The contract should connect requirements to the delivered service and responsible party.

Customer data can pass through storage, memory, logs, support systems and telemetry. The data map identifies location, purpose, retention, access and transfer. Training or model-improvement use requires an explicit agreed basis. Deletion should be technically and contractually defined across primary, backup and support records.

Continuity planning should consider utility outage, cooling failure, network loss, hardware fault, cyber incident, supplier failure, civil emergency and customer transition. The recovery design should match the contracted product. A single facility cannot promise geographic recovery without available capacity, data and tested procedures elsewhere.

Incident communications need severity, clock, audience, content, update frequency and authority. A premature or inaccurate statement can create further harm. The operating process should preserve facts, legal requirements and customer needs.

Audit rights should be proportionate to multi-tenant security. Independent reports, certifications, targeted evidence and controlled visits can provide assurance. Highly sensitive system information requires protection.

19. Operate a 120-day contract and financing office

During days one to twenty, management defines the target customer, workload, capacity product, jurisdiction, service boundary and financing objective. It reconciles site, utility, facility, equipment, cooling, network and operating evidence. Material gaps enter a controlled register.

During days twenty-one to fifty, the team builds the capacity schedule, tariff model, acceptance plan, service metrics, credit package, technology roadmap and legal term sheet. Customer and provider assumptions remain separate. Engineering, commercial, legal, finance and operations owners approve their boundaries.

During days fifty-one to eighty, the provider conducts customer diligence, technical workshops and contract negotiation. It obtains indicative or committed supplier terms, validates power and cooling capability, maps network routes and tests the illustrative workload. The financial model is updated with accepted evidence.

During days eighty-one to one hundred, the parties agree the acceptance pack, reporting, billing, security, incident, change, refresh, remedy and transition schedules. Financiers review the contract and asset evidence. Open conditions, consents and direct agreements are assigned.

During days one hundred and one to one hundred and twenty, the board approves the final capacity allocation, downside case, customer exposure, capital commitment and signing authority. The implementation office baselines milestones and evidence. No signature should precede required approvals.

Table 6. 120-day AI capacity contract office

Period	Decision output	Core evidence	Commercial gate
days 1-20	defined capacity product	demand, site, power, cooling, network and equipment map	credible service boundary
days 21-50	tariff and term-sheet architecture	cost model, metrics, acceptance and credit design	priced controllable risk
days 51-80	diligenced customer and solution	credit, engineering, supplier and workload evidence	deliverable customer case
days 81-100	executable contract schedules	billing, security, remedies, refresh and transition	complete documentation
days 101-120	board and financing decision	final model, approvals, conditions and direct rights	funded implementation
recurring	operating contract control	telemetry, invoices, incidents, credits and capacity	retained oversight

Timing is an illustrative management assumption; the actual critical path follows project evidence and third-party decisions.

20. Make the investment decision on reconciled evidence

The board should approve one capacity definition, one customer obligation, one engineered delivery plan, one tariff model and one downside case. It should understand which revenue is reserved, accepted, consumed, billed and collected. Each stage has a different evidential status.

The decision paper should state the customer credit exposure, capacity concentration, equipment commitment, power and thermal headroom, network dependency, acceptance critical path, service remedy, refresh requirement, termination exposure, refinancing need and replacement strategy. Material assumptions should remain visible.

The commercial opportunity extends beyond drafting a contract. Sponsors can need capacity-product design, customer segmentation, pricing architecture, term-sheet support, customer diligence, supplier coordination, financing materials, data-room construction, acceptance readiness and recurring contract oversight. These are defined strategy-and-execution workstreams with accountable deliverables.

Benefits remain prospective until demonstrated. A well-designed tariff can improve comparability, capital allocation and operating control. A customer can receive a clearer service and remedy system. A financier can obtain a more coherent evidence bridge. Each result depends on implementation and the executed arrangement.

The central conclusion is that megawatts, accelerators, cooling and network become financeable demand only when the contract turns them into one measurable, accepted and collectible customer service. The tariff prices that service across reservation, use, performance and change. The board then decides on evidence rather than headline capacity.

References

[1] - International Energy Agency. Energy and AI, 10 April 2025. https://www.iea.org/reports/energy-and-ai

[2] - International Energy Agency. Energy Demand from AI, Energy and AI. https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai

[3] - International Energy Agency. Data Centre Electricity Use Surged in 2025, 16 April 2026. https://www.iea.org/news/data-centre-electricity-use-surged-in-2025-even-with-tightening-bottlenecks-driving-a-scramble-for-solutions

[4] - United States Department of Energy. Powering America's AI Future: Data Center Resource Hub. https://www.energy.gov/powering-americas-ai-future-data-center-resource-hub

[5] - United States Department of Energy. DOE Releases New Report Evaluating Increase in Electricity Demand from Data Centers, 20 December 2024. https://www.energy.gov/articles/doe-releases-new-report-evaluating-increase-electricity-demand-data-centers

[6] - United States 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

[7] - United States Department of Energy. Best Practices Guide for Energy-Efficient Data Center Design, 26 July 2024. https://www.energy.gov/cmei/femp/articles/best-practices-guide-energy-efficient-data-center-design

[8] - United States Department of Energy. Cooling Water Efficiency Opportunities for Federal Data Centers. https://www.energy.gov/cmei/femp/cooling-water-efficiency-opportunities-federal-data-centers

[9] - European Commission. Commission Delegated Regulation (EU) 2024/1364 of 14 March 2024. https://eur-lex.europa.eu/eli/reg_del/2024/1364/oj

[10] - National Institute of Standards and Technology. Cloud Computing Service Metrics Description, NIST SP 500-307, 2018. https://doi.org/10.6028/NIST.SP.500-307

[11] - National Institute of Standards and Technology. Service-Level Agreement Glossary. https://csrc.nist.gov/glossary/term/sla

[12] - Open Compute Project Foundation. Data Center Liquid Distribution Guidance and Reference Designs, 2026. https://www.opencompute.org/documents/ocp-acf-reference-design-guidance-white-paper-pdf-1

[13] - Open Compute Project Foundation and ASHRAE. New Alliance Focused on Data Center Liquid Cooling Technology Standards and Best Practices, 13 October 2025. https://www.opencompute.org/blog/open-compute-project-foundation-and-ashrae-form-new-alliance

About the Author

Chennakeshav Adya, Independent Researcher