



Paul Turner
Turner Commodities Law, PLLC
713-899-5027
paul.turner@turnercommoditieslaw.com

James Guy
Deacon Law Group, PLLC
512-576-2435
jamesguy@deaconlawgroup.com



Powering the Hyperscaler

*A Legal and Commercial Guide to Gas, Power, Trading and Regulatory Risk
in Serving Data Centers in the Electric Reliability Council of Texas*

By Paul Turner, Principal
Turner Commodities Law, PLLC
James Guy, Principal
Deacon Law Group, PLLC
October 7, 2026

Introduction

The current wave of hyperscale data centers is driven by artificial intelligence. Training and running AI models requires dense concentrations of computing power, and therefore unusually large amounts of electricity — far more per square foot than conventional computing.

Computing equipment converts electricity into heat that must be removed. Cooling connects the data-center boom to another constrained resource: water. Water can be consumed directly in cooling systems and indirectly in generating the electricity used by the facility. The amount varies with cooling technology, power source, location and operating profile, but Texas regulators and policymakers are increasingly examining computing, electricity and water together.

Until recently, Texas did not regulate data centers through a single, integrated regulatory framework. A large facility principally encountered the same bodies of law that applied to other unusually large industrial loads: utility interconnection requirements, Electric Reliability Council of Texas (ERCOT) reliability and market rules, retail and wholesale electric regulation, transmission planning, environmental permitting, water law and ordinary commercial contracting.

That system developed when very large new loads were comparatively unusual. Before 2022, a large customer seeking service generally worked directly with the transmission or distribution service provider in whose territory it planned to locate. ERCOT did not have a single, system-wide process for coordinating the effects of numerous very large loads seeking to connect at roughly the same time. In 2022, ERCOT introduced an interim large-load process and later developed a formal Large Load Interconnection Study process.

The enormous influx of data-center projects exposed a weakness in the sequential study model. Each project was studied against the system as it then existed. If another large project affected the same transmission facilities, previous study results could become outdated and projects that had not yet energized could be required to undergo additional study. ERCOT has described these repeated re-study cycles as potentially extending project timelines by years.

This result was more than a regulatory inconvenience. Interconnection uncertainty became a commercial problem. A data-center developer might be negotiating a power arrangement while waiting for transmission studies. A generator might be reserving turbines or construction capacity. A marketer might be considering long-term gas supply or hedges. A pipeline might be asked to reserve or expand transportation capacity. Yet none of those commitments necessarily had the same development milestones, termination rights or timing assumptions.

Texas policymakers have begun changing that framework. Senate Bill 6 directed the Public Utility Commission of Texas (PUCT) to establish more rigorous standards governing large-load interconnection, forecasting, financial commitment and related reliability matters. ERCOT has also moved toward coordinated batch analysis for large-load projects. For energy companies considering whether to serve hyperscale demand, the legal and regulatory scheme answers only part of the question.



Paul Turner
Turner Commodities Law, PLLC
713-899-5027
paul.turner@turnercommoditieslaw.com

James Guy
Deacon Law Group, PLLC
512-576-2435
jamesguy@deaconlawgroup.com



DEACON LAW GROUP

The more fundamental commercial issue is that a hyperscale data-center transaction is rarely just a power purchase agreement. Depending on the structure, serving one large campus may require an interlocking series of arrangements involving natural gas supply, pipeline transportation and possibly storage, dedicated generation or tolling, wholesale power, retail electric service, qualified scheduling entity (QSE) representation, ERCOT scheduling and settlement, transmission and congestion, financial hedging, demand response and curtailment, water, environmental permits, interconnection infrastructure, and substantial credit and development support.

Different entities may perform each of those functions. Different agreements regularly govern them. And different regulatory regimes may apply to them. That complexity creates a deceptively simple legal problem: the transaction must work not only contract by contract, but also as a system.

A power agreement can be carefully drafted and still leave the supplier exposed if the gas transportation agreement continues imposing reservation charges while the customer has a right to delay commercial operation. A tolling arrangement can allocate fuel responsibility clearly while leaving uncertainty over whether the generator or the QSE controls dispatch. A hyperscaler may have extraordinary corporate credit but still retain the right not to proceed with a particular campus after its energy counterparty has committed millions of dollars to equipment, infrastructure or hedges.

Similarly, a twenty-year power arrangement may create exposures that cannot be economically hedged for twenty years. A data center may agree to curtail load during stressed ERCOT conditions while the generator remains committed to take gas. A change in ERCOT's transmission-cost recovery methodology may be passed through under one agreement but absorbed under another. The greatest risks therefore often appear in the seams between the contracts.

For trading companies, generators, marketers and integrated energy companies, hyperscale development presents a potentially significant new source of long-duration demand for both natural gas and electricity. It can support new generation, pipelines, structured commodity transactions, hedging and optimization. But those opportunities also can require the energy provider to commit capital, credit and market positions well before the customer's load becomes operational.

This article examines the questions trading, commercial, legal, regulatory, risk and compliance personnel should ask before committing capital or balance sheet to serve hyperscale load.

Executive Summary

Hyperscale data centers are creating a potentially important new market for natural gas, generation, and electricity in Texas. For commodity companies, capturing that opportunity requires looking beyond the power contract.

A large campus creates a chain of interdependent obligations: gas procurement supports generation, pipeline transportation supports gas procurement, generation supports a power obligation, a retail electric provider (REP) or bundled utility serves the load, a QSE interfaces with ERCOT, and financial hedges may be layered across several of those exposures. The project itself may depend on interconnection, transmission upgrades, water, environmental permits, equipment availability, construction, financing, and continued commitment by the hyperscaler.

The legal task is to make those pieces fit together — and to test the transaction against the downside case, not just the model. The greatest risks often appear in the seams between the contracts: a power agreement that permits delay while the gas transportation agreement continues billing reservation charges, a tolling arrangement that leaves dispatch control ambiguous, a hedge that settles after the customer reduces its obligation.

The central question is not simply whether the supplier can deliver the power. It is what the supplier has committed across the entire gas, power, transportation, market, and regulatory structure — and what happens to those commitments if the project develops differently from the base case.



Paul Turner
Turner Commodities Law, PLLC
713-899-5027
paul.turner@turnercommoditieslaw.com

James Guy
Deacon Law Group, PLLC
512-576-2435
jamesguy@deaconlawgroup.com



I. Why Data Centers Are Different

A data center houses the computers that store information and run digital services. A hyperscale facility does the same work on a much larger scale, often across multiple buildings on a single campus. Its servers need electricity, but so do the equipment that moves information through the facility, the systems that convert and distribute power, and the cooling systems that remove heat. The relevant commercial load is therefore the facility as a whole, not merely the nameplate rating of its computing equipment.

These facilities also place a high value on reliable service. Computing equipment, communications and cooling must operate together, and an interruption can disrupt the services the data center provides. Operators may use batteries, backup generation, redundant equipment and more than one path for electric service to manage that risk. However, those arrangements do not, by themselves, tell an energy supplier how much power the customer will buy at each hour or what the customer can reduce during an ERCOT emergency. The answers depend on the facility's design and the rights the parties put in their agreements.

Serving a hyperscale campus can require new transmission, dedicated generation, natural gas supply and transportation, retail service, ERCOT market representation and water infrastructure. Each element may follow a different schedule and contract.

II. From an Industrial Load to a New Regulatory Category

For many years, a proposed Texas data center approached electric service much as another large industrial customer would. It worked with the transmission or distribution utility authorized to serve its site, while ERCOT studied the effect of adding the load to the wider integrated system. That approach became harder to manage as requests for very large amounts of capacity at different locations arrived at the same time. One project could alter the assumptions used to study another, requiring further study before either could confidently set an energization date. Further, as regulatory processes slowed and the race for electric service continued, data-center developers began submitting requests for the same projects at different locations, leading to uncertainty about the magnitude of very large loads in the regulatory pipeline. As a result, the base assumptions for large load interconnection studies became increasingly difficult to isolate.

Texas has responded by treating very large loads as a distinct interconnection problem. Senate Bill 6 directed the PUCT to establish standards for large loads in the ERCOT region that support development while limiting stranded infrastructure costs and protecting reliability. The statute addresses site control, disclosure of substantially similar service requests, study fees, financial commitments, load forecasting, and information about certain onsite backup generation. It sets 75 MW—whether new or expanded to at a single customer's site—as the presumptive threshold for the new standards.

The PUCT took a further step in September 2026 by adopting its large-load interconnection rule. As adopted, the rule requires a \$100,000 initial study fee and financial security tied to the amount of capacity requested—\$50,000-per-MW, though the financial security is potentially refundable. The adopted rule also lengthened the period before unused capacity must be reported to ERCOT: the final rule uses a 24-month period measured against the energization schedule as a whole, not against each milestone individually (as the original proposal had done).

The adopted rule operates alongside ERCOT's related batch study process and a separate verification pause. On August 3, 2026, Governor Abbott directed the PUCT and ERCOT to audit data centers advancing through ERCOT's interconnection process before additional data centers are approved to move forward. At the time, the Governor's office reported approximately 474 GW of connection requests, roughly 90 percent attributable to data centers—more than five times ERCOT's record peak demand. ERCOT paused the Batch Zero study process and delayed final large-load classifications the same day, then requested regulatory relief on August 10. ERCOT's announced schedule targets December 10, 2026, for filing its Batch Zero Eligibility Verification Report; that aspiration is not a guaranteed date for final classifications or resumption or completion of the batch study. The pause does not repeal the adopted



Paul Turner
Turner Commodities Law, PLLC
713-899-5027
paul.turner@turnercommoditieslaw.com

James Guy
Deacon Law Group, PLLC
512-576-2435
jamesguy@deaconlawgroup.com



rule: its financial commitments and study fees apply within a framework whose qualifying loads and study timetable remain under review.

Site control, a credible energization schedule, study payments and financial security provide evidence of a developer's commitment. They do not establish that the full requested load will be built, ramp on schedule, or produce payments to an energy supplier for capacity reserved in anticipation of it, but they represent a preliminary hurdle cleared.

That distinction matters when the supplier must act before the data center is energized. A generator may need to order equipment; a gas marketer may need to acquire supply; and a pipeline customer may need to reserve firm transportation. Each commitment has its own date and its own cost if the campus is delayed, reduced, or abandoned. Then, there are upgrades necessary for the ERCOT grid to support the interconnection. Transmission, once built, becomes incorporated into the cost of providing transmission service for the entirety of the ERCOT grid regardless of delayed or abandoned projects. Interconnection security protects interests defined by the interconnection framework. It does not automatically reimburse the energy company for those separate commercial exposures, nor does it fully protect other electric customers from the costs of serving data centers.

The diligence question is therefore more precise than whether the project is "in the queue." What capacity has been requested, studied, secured, contracted for, and actually made subject to a binding payment obligation? At each stage, the parties should ask who pays if the data center's schedule slips while the gas, generation, transportation, or hedge commitments remain in force.

III. Think of the Transaction as a Contract Stack

A data center may describe its need simply: reliable power at an agreed price. Meeting that need can require several companies to perform different jobs under separate contracts. The REP or utility serves the customer. A marketer may buy power or manage price risk. A QSE communicates with ERCOT and handles market scheduling and settlement. A generator may produce some or all of the electricity. If that generator burns natural gas, still other agreements may govern fuel supply and pipeline transportation. Hedging programs may provide additional risk protection.

The companies do not necessarily contract with one another in a straight line. A hyperscaler might contract for delivered power from a marketer, while a separate affiliate develops generation near the campus, while both entities might contract with the REP or electric utility to provide retail service. The marketer might obtain power under a tolling agreement and arrange gas supply but rely on another company for transportation and a QSE for ERCOT functions. The data center could also retain grid service, backup generation, batteries, or the ability to reduce its load.

Each contract can be reasonable on its own terms while the combined transaction leaves a substantial risk unassigned. The customer's power agreement might permit a year of delay, for example, while the generator must begin paying for firm gas transportation when a new pipeline enters service. A tolling agreement might give the customer dispatch rights without giving the gas supplier enough notice to change its nominations. A hedge might continue settling after the customer reduces the amount of power it is required to buy. None of those problems is necessarily visible from a review of the power agreement alone.

The same issue arises with credit support. A parental guarantee may secure the customer's payment obligations but exclude development costs incurred before service begins. An interconnection deposit may protect a utility against specified costs without protecting a generator or marketer that has ordered equipment or entered into a long-term fuel commitment. Counsel should identify the party entitled to each form of security and the obligations it actually covers. As discussed in Section II, the regulatory framework also imposes study-payment and financial-security requirements. Those requirements serve their specified regulatory purposes and should be distinguished from the commercial credit support needed to protect the energy supplier's commitments.



Paul Turner
Turner Commodities Law, PLLC
713-899-5027
paul.turner@turnercommoditieslaw.com

James Guy
Deacon Law Group, PLLC
512-576-2435
jamesguy@deaconlawgroup.com



For each agreement, identify who performs, who controls operational decisions, when obligations begin, what events excuse performance, what costs continue during an interruption, what regulatory and compliance obligations each party assumes, and what happens if the data center is delayed or built at a smaller scale. Then compare the answers across agreements. If one party can change the project's schedule or operating profile, the documents should say who bears the resulting costs elsewhere in the structure.

IV. Start with the Load: Commercial and Development Diligence

A proposed data-center campus is often described by its ultimate capacity: 500 MW, 1 GW, or more. That figure may be useful for planning, but it is a poor substitute for the amount of load an energy supplier can count on serving. A project can have an announced size, an interconnection request, a studied amount, an approved service arrangement, a contracted purchase obligation, and an actual operating load. Those numbers may differ substantially. Counsel should ask what each figure means, who supplied it, and what event would make it change.

Start with the physical project. Does the developer control the site for the full proposed buildout? What transmission facilities and upgrades are needed, and where is the project in the interconnection process? Have the buildings, cooling systems, water supply, backup generation, and any dedicated power plant been designed and permitted? Which equipment has been ordered, and which construction dates depend on equipment that is difficult to obtain? A schedule supplied to a utility is relevant evidence, but it is not a substitute for checking the development work behind it.

Then examine who is committed to the campus. A developer may have a credible site and interconnection request but no final decision by the hyperscaler (or multiple smaller tenants) that will occupy or use the facility. The agreement may cover an initial building while leaving later buildings to the customer's option. Financing, tenant commitments, board approvals, and rights to relocate computing demand can all affect whether the projected load appears at that particular site. Evaluation of the campus should identify the decisions that remain open when the energy provider is asked to commit.

Phasing deserves particular attention. A campus described as 1 GW might begin with a much smaller first phase and add capacity over several years. The parties should specify the minimum load and ramp for each phase, the conditions for moving to the next one, and whether the customer may defer or decline a phase without paying for energy infrastructure built in anticipation of it. A right to delay the final 750 MW is commercially different from a binding obligation to take it on a stated date, even if both projects appear in planning materials as a 1-GW campus.

The critical contract question is when expected demand becomes an enforceable obligation. A commitment to request interconnection, post security, or use reasonable efforts to build does not necessarily require the customer to buy electricity. Conversely, a minimum payment obligation can support the supplier's investment even if actual consumption is lower than expected. The power agreement should state the capacity and payment obligations at each development milestone, the permitted extensions and reductions, and the amount payable if the customer stops the project. Any parental guarantee should cover the obligations that matter before and after commercial operation.

These dates should be compared with the supplier's own commitments. When must it order turbines, reserve pipeline capacity, construct a lateral distribution facility, procure power, or begin hedging? Can it postpone those steps until the customer's obligation is firm? If it cannot, the parties should address reimbursement, development security, termination payments, or the supplier's right to use or remarket stranded assets. Creditworthy customers can pay what they owe; the document must first make clear what they owe if the project is late or never reaches its announced size.

The result should be a staged view of the load, not a single headline number. For each phase, the approval team should be able to identify the capacity requested, the capacity that can be energized, the capacity the customer must pay for, the expected operating ramp, and the supplier commitments already made.



Paul Turner
Turner Commodities Law, PLLC
713-899-5027
paul.turner@turnercommoditieslaw.com

James Guy
Deacon Law Group, PLLC
512-576-2435
jamesguy@deaconlawgroup.com



V. Gas Supply and Transportation

If a data center will rely on gas-fired generation, the power arrangement depends on a second supply chain. Someone must buy gas in the right quantities, move it to the plant, and manage the difference between scheduled deliveries and actual burn. Those tasks can be divided among the generator, a gas marketer, a tolling customer, and one or more pipeline shippers. The contracts should identify who does each task before they allocate the resulting costs.

The starting point is the plant's expected operating pattern, not merely its maximum heat rate multiplied by its nameplate capacity. A generator serving a steady campus load may still start, stop, ramp, or run below its expected output because of maintenance, changing load, ERCOT conditions, or the economics of buying power instead. Fuel diligence should test hourly and daily burn, start-up fuel, minimum operating levels, cycling limits, gas quality and pressure, and the consequences of a rapid change in dispatch. The gas purchase terms must be workable for that operating pattern, including the index, delivery point, quantity flexibility, nomination deadlines, and any obligation to take or pay for unused gas.

Transportation is a separate decision. Firm pipeline capacity generally offers higher-priority service in exchange for a reservation charge that may be due even when little or no gas moves. Interruptible service may cost less but can disappear when capacity is scarce. The appropriate mix depends on how the power obligation treats fuel interruption and replacement power. A promise to deliver firm electricity is difficult to price if the generator assumes it can always obtain interruptible gas transportation during the hours when the grid is under stress.

The review should identify the pipeline receipt and delivery points, maximum daily quantity, upstream and downstream interconnections, any lateral or compression that must be built, and whether the shipper can use alternate points or release unused capacity. A gas supply contract at one hub does not solve a constraint on the path to the plant. Nor should the team assume that capacity release will recover the full reservation charge: that depends on the applicable pipeline arrangements, available buyers, and the value of the capacity when it is offered.

Who submits nominations, receives pipeline notices, changes schedules after a dispatch instruction, and communicates with the plant and QSE? Who pays for imbalances, cash-outs, fuel retention, unauthorized overruns, and penalties under an operational flow order (OFO)? A late change in electric dispatch may be economically sensible in ERCOT while leaving the gas shipper unable to change its scheduled deliveries without cost. The party controlling dispatch should see those costs, and the documents should say who bears them.

Development timing presents the clearest example of the contract stack problem. Suppose a marketer commits to fifteen years of firm transportation to support dedicated generation. The pipeline and plant are ready, but the data center postpones its first phase for two years under the power agreement. The pipeline may perform exactly as promised and continue billing reservation charges throughout the delay. The power and gas documents should specify whether the customer reimburses those charges, whether the marketer can use or remarket the capacity, how any proceeds are credited, and what happens if the delayed phase is never built.

If the campus reduces load, a unit is on outage, or the customer exercises a curtailment right, gas purchases and transportation charges may continue under their own terms. The transaction team should test those cases against the power price, minimum payments, fuel-supply flexibility, transportation rights, and any hedge. The objective is a practical answer to a simple question: when the plant does not burn the gas originally expected, who owns the gas, the pipeline capacity, and the cost?

VI. Generation, Power Marketing and Tolling

In the ERCOT market, a large data center can be served from a supplier's generation portfolio, a plant developed chiefly for that campus, or a combination of dedicated generation and market purchases. The choice changes the risks the supplier is taking. A portfolio can replace the output of one unavailable unit with other resources or market purchases, although the cost may be high. A dedicated plant gives the customer a clearer connection to identified



Paul Turner
Turner Commodities Law, PLLC
713-899-5027
paul.turner@turnercommoditieslaw.com

James Guy
Deacon Law Group, PLLC
512-576-2435
jamesguy@deaconlawgroup.com



capacity but concentrates construction, outage, fuel, and performance risk in that facility. The power contract should say whether the customer is buying electricity delivered at a point or rights associated with particular generating units.

A tolling arrangement makes that distinction especially important. Under a typical toll, the customer or toller provides or pays for fuel and receives rights to the plant's electric output in exchange for capacity and other payments. But the label "tolling" does not answer who decides when a unit starts, how it is offered into ERCOT, or what happens if it consumes more fuel than the assumed heat rate. The agreement must specify the dispatch rights, operating limits, fuel delivery responsibilities, start and shut-down costs, minimum run times, ramp limits, and treatment of tests and commissioning. It should also distinguish the right to make an economic dispatch decision from the plant operator's responsibility to operate the equipment safely and comply with applicable instructions.

The parties should trace every dispatch decision to its market and commodity consequences. If the toller requests a start, who submits the required information through the QSE, buys the gas, changes pipeline nominations, and pays for start fuel? If ERCOT directs a change in output, who receives the instruction and tells the plant and gas shipper? If the toller chooses not to run when the data center needs power, who buys replacement energy and bears congestion or price spikes? Those answers may be spread among the toll, the QSE agreement, the gas arrangements, and the customer's power agreement. They should be consistent and not result in a party committing more financial security than is needed.

A plant may meet an annual availability target yet be unavailable during the few hours when replacement power is most expensive. The documents should address scheduled maintenance, forced outages, notice, testing, derates, and whether other customers or generating units may receive priority. They should also state whether the supplier owes replacement power, a payment credit, damages, or some combination when the plant cannot perform. A cap on one remedy should be compared with the supplier's uncapped obligation, if any, to serve the data center under a separate agreement.

Construction and fuel assumptions can move these risks before commercial operation. Turbine delivery, interconnection, gas pipeline facilities, air permits, water supply, and commissioning may each follow a different schedule. If the data center is ready before the plant, the supplier needs a bridge supply plan and a price for it. If the plant is ready first, someone must pay its fixed costs and any committed transportation while the customer's load is absent. The agreements should define the relevant commercial-operation dates separately and explain what happens when they do not coincide.

Environmental and equipment limits also constrain the dispatch right being sold. Permit limits on emissions or operating hours, manufacturer requirements, long-term service agreements, and maintenance cycles may prevent a plant from following every economically attractive instruction. A contract that grants broad dispatch control should identify those limits, establish who maintains compliance, and allocate the cost of using up a limited operating allowance or accelerating maintenance. The party with the economic benefit of a dispatch decision may be different from the party legally responsible for the unit.

The appropriate question is therefore not simply whether the plant can produce the required MW. It is who may direct it, who communicates with ERCOT, who supplies and schedules fuel, and who pays when physical performance differs from the power obligation. Those responsibilities should remain clear whether the energy comes from a named plant, the supplier's portfolio, or the wholesale market.

VII. Retail Service, Scheduling and ERCOT

A data center power arrangement needs more than a seller and a buyer. In a Texas area open to retail competition, a REP sells electricity to the customer and handles the retail relationship. In an area not open to retail competition, the retail provider might be an electric cooperative or a municipally owned utility (MOU). In both areas, a QSE represents market participants in ERCOT, including load-serving entities such as REPs, cooperatives, MOUs and entities



Paul Turner
Turner Commodities Law, PLLC
713-899-5027
paul.turner@turnercommoditieslaw.com

James Guy
Deacon Law Group, PLLC
512-576-2435
jamesguy@deaconlawgroup.com



responsible for generation or qualified load resources. The customer, retail provider, generator, and QSE may be different companies. Their contracts should identify each company and the job it has agreed to perform.

The retail provider remains the retail-facing party even if another company supplies the generation or performs scheduling and billing work for it. The retail service agreement should say who obtains meter data, forecasts the campus load, procures energy and related services, bills the customer, and explains adjustments after ERCOT settlements change. It should also identify who pays transmission and distribution charges, ERCOT-related charges, and taxes, and whether each is included in the stated power price or passed through. Delegating a task to an affiliate or contractor should not leave the customer uncertain about whom to call or which company stands behind the retail promise.

The QSE agreement answers a different set of questions. It should define the QSE's authority to submit market bids, offers, schedules, and resource information; the data and instructions it receives; the deadlines for those instructions; and the process for checking ERCOT statements and correcting errors. If a dedicated plant and the data center load have different QSEs, the parties need a reliable way to coordinate their forecasts, dispatch decisions, outages, and settlement data. They should identify the person or control room authorized to speak for each side around the clock. A commercial instruction to run a plant or reduce load must reach the entity able to implement it in time.

A data center's physical operating choices can also affect its market position. A sudden step up in consumption, a transfer to backup generation, or a shutdown may leave the retail provider or QSE buying or selling a different quantity than expected. In addition to changing the quantity, this can create a change in regulatory and market compliance obligations such as credit exposure for the QSE. The agreements should specify how promptly the customer must give notice of commissioning, phased energization, planned maintenance, material changes in load, and unplanned interruptions. They should allocate the resulting imbalance, congestion, and settlement exposure without treating every deviation alike: a customer-directed change, an ERCOT instruction, and an outage at the supplier's plant have different causes.

The right to offer flexible load into an ERCOT program needs an express owner. A qualified load resource can participate through a QSE, but enrollment, telemetry, testing, availability commitments, performance, and compensation require coordination with the customer's operations. The retail and QSE agreements should state who may enroll the load, who can commit a quantity, who receives program revenue, and who pays if the load fails to respond. They should also prevent two parties from selling the same flexibility or assuming that a campus can curtail when its operating requirements do not allow it. The commercial treatment of curtailment warrants its own discussion below.

Finally, the parties should map failure and transition. If the retail service company, QSE, or generator changes, how will customer enrollment, market representation, telemetry, data access, and outstanding settlement disputes transfer? If ERCOT later corrects a settlement for a period before the change, who receives the charge or credit and who may challenge it? A short responsibility schedule attached to the agreements can make these answers visible.

VIII. Hedging, Marketing and the Twenty-Year Problem

A data center may want a predictable power price for fifteen or twenty years. The supplier, however, may have to buy gas, power, transmission-related products, and financial hedges on much shorter terms. Signing the customer contract does not lock in the cost of serving it for the full term. The unhedged years, and the dates when existing hedges must be replaced, are part of the bargain even if they do not appear on the first-year price sheet.

The exposure is more complicated than the direction of wholesale power prices. Gas bought at one location may need to reach a plant at another. A power hedge may settle at a trading hub while the plant injects at its node and the data center takes service in a load zone. Congestion can change the difference between those prices; a congestion revenue right can help address a defined path, but it does not make every location or hour identical. A flat block of



Paul Turner
Turner Commodities Law, PLLC
713-899-5027
paul.turner@turnercommoditieslaw.com

James Guy
Deacon Law Group, PLLC
512-576-2435
jamesguy@deaconlawgroup.com



DEACON LAW GROUP

power may also be a poor match for a campus whose load grows in stages, varies with cooling demand, or transfers to backup generation. Each mismatch can leave a residual cost after the headline hedge performs as written.

A gas-fired plant adds another moving part. The expected relationship between gas cost and power value depends on the unit's heat rate, availability, operating limits, and the hours it runs. A fixed gas position can become too large when the plant is down or the customer delays taking load; it can become too small when output or gas consumption exceeds the assumed profile. Someone then has to sell excess gas, buy replacement gas, or buy power while the underlying customer price may be fixed. The toll, gas supply, hedge documents, and retail power agreement should allocate those outcomes consistently.

The twenty-year problem is that a long physical obligation may outlast the useful or available tenor of the supplier's hedges. Even where a forward product is listed for future years, the quantity, shape, location, and price available for a large project may not match its full obligation. A rolling hedge program requires future trades at future prices and credit terms that cannot be known when the customer contract is signed. The supplier should model what happens at each hedge roll and under stressed prices, rather than treating the current forward curve as a guaranteed procurement plan.

Credit risk runs in both directions. A creditworthy data center customer may pay its invoices on time while falling prices create a large mark-to-market loss on the supplier's hedges, or rising prices cause a margin call before the customer pays for delivered power. If the project is delayed, reduced, or terminated, the supplier may have to unwind positions at a loss even if no electricity was delivered. The agreements should address security, margin and liquidity needs, the timing of customer payments, early termination amounts, and whether a parent guarantee or other support actually covers the exposure at the entity that bears it. A cap on the customer's liability should be tested against the supplier's possible hedge unwind and replacement costs.

The commercial team can reduce some of this risk through the price design. Possible terms include a shorter fixed-price period within a longer service term, indexed components for gas or power, a stated mechanism for resetting prices when hedges expire, volume bands, separate pricing for load growth, and an agreed treatment of congestion or extraordinary charges. Each choice shifts a specific risk to the customer or supplier; it should be described in the offer rather than buried in a pass-through definition. The supplier also needs trading limits and authority to decide when, how much, and for whom to hedge.

Before approving a long-term offer, decision-makers should see the expected load by year and hour, the committed generation and fuel, the hedge coverage by product and tenor, the remaining open positions, and the liquidity required in adverse scenarios. The question is whether the business can carry the risks that remain after the contracts and hedges are read together. A strong customer covenant supports payment for service; it does not itself supply a twenty-year hedge or cash for a margin call tomorrow.

IX. Curtailment and Flexible Load

A data center's ability to reduce grid demand can have value, but "flexible load" is not a single product. A campus might voluntarily reduce computing work when power prices rise, commit a measured reduction through an ERCOT program, switch equipment to backup generation, or be required to reduce load during a potential grid emergency. Each has a different trigger, notice period, operating limit, and economic result. The customer and supplier should specify which forms of flexibility the project can actually deliver before assigning anyone the right to sell them.

Texas law now adds a distinct emergency obligation for certain large loads. Senate Bill 6 calls for a curtailment protocol for qualifying transmission-voltage load interconnected after December 31, 2025, during firm load shed, subject to the statute's stated exceptions. It also directs development of a service to competitively procure demand reductions from customers with demand of at least 75 MW for anticipated emergencies. That service requires at least 24 hours' notice and excludes a customer that curtails in response to wholesale prices or participates in another reliability or ancillary service. The PUCT's implementing proposal in Project No. 58482 was still under consideration



Paul Turner
Turner Commodities Law, PLLC
713-899-5027
paul.turner@turnercommoditieslaw.com

James Guy
Deacon Law Group, PLLC
512-576-2435
jamesguy@deaconlawgroup.com



DEACON LAW GROUP

as of this draft. A contract should preserve room to comply with the final requirements rather than assume that every form of curtailment may be stacked or paid twice.

Separately, qualified load resources can participate in ERCOT energy and ancillary-service markets through a QSE. A customer may also respond economically to its own power price without enrolling in a program. The contract should distinguish a mandatory instruction from ERCOT or the utility, a market commitment made in advance, and a customer choice made for its own account. It should say who may offer the load, who receives an instruction, who decides whether the required reduction is technically possible, and who bears the consequences of failing to perform a commitment, including any regulatory and market compliance consequences. The retail service provider, QSE, operator, and generator need a common view of those decisions.

A campus may be able to defer a batch of computing work but still need electricity for cooling, communications, security, and a controlled shutdown. The parties should define a minimum protected load, achievable power reductions, response time, duration, frequency, recovery ramp, and any limits on shifting work between buildings or sites. Backup generation cannot be treated as unlimited flexibility: fuel inventory, air permits, maintenance, and the time needed to transfer load can constrain its use. The utility's curtailment protocol and the customer's internal operating procedure should be tested together.

If the load drops while a dedicated plant continues to run, who markets the excess output and receives the proceeds? If the plant backs down, who changes gas nominations and pays for minimum volumes or unused transportation? If the supplier bought power or hedged the expected load, who receives gains or pays unwind losses? A reduction in the customer's meter reading does not by itself cancel a capacity payment, a fuel commitment, or a financial hedge. The power agreement, toll, gas contracts, and QSE arrangements should use consistent settlement rules for the same event.

The value of flexibility should be accounted for once. Program payments, avoided energy purchases, released generation, and changes in transmission charges can accrue to different parties under different contracts. The documents should state what the customer keeps, what the supplier keeps, and whether a credit against the power invoice applies. They should identify the baseline and meter data used to measure performance, the treatment of partial reductions, and a process for correcting later ERCOT settlement changes. The same MW should not be promised to two programs or priced twice in the customer's deal.

A useful approval test is to walk through a specific event: ERCOT gives notice of a possible emergency tomorrow; the data center reduces by 100 MW for four hours; its plant remains available; and gas has already been nominated. The parties should be able to identify who receives the notice, who directs the campus and plant, what can safely be reduced, how ERCOT sees the response, and where every resulting cost or payment lands. If those answers are still scattered among contracts, the project's claimed flexibility may be less valuable than it looks.

X. Water, Air and Other Non-Power Conditions

The principal focus of this article is on energy-related agreements, rather than the water, air, land and other arrangements needed to develop and operate a data center. Similar approval hurdles and risks at the seams between agreements nevertheless arise in those areas. A power supply can be ready before the campus has secured a dependable water source, completed its cooling equipment, arranged wastewater service or obtained permits for backup engines. A dedicated plant may face separate water and air requirements. If one of those conditions delays the campus, generation, gas transportation, reserved transmission capacity and hedges may already be costing money.

For purposes of the energy transaction, diligence should establish how the water source, rights, infrastructure and available volume support the campus's planned growth, and how drought or an outage could affect operations. Cooling choices can alter both water demand and the power load used to price the deal. Air authorizations should be checked against the intended operation of the dedicated plant and backup engines; an engine approved for emergency use should not be assumed available for routine economic dispatch. The parties should coordinate



Paul Turner
Turner Commodities Law, PLLC
713-899-5027
paul.turner@turnercommoditieslaw.com

James Guy
Deacon Law Group, PLLC
512-576-2435
jamesguy@deaconlawgroup.com



construction, wastewater and other site approvals with interconnection and plant completion, identifying who controls each prerequisite and which dependencies could delay energy service.

The same contractual discipline applies across these arrangements. The documents should identify which approvals are conditions to construction, first energization and full service, and allocate the consequences if an approval is late, denied or materially restricted. Long-stop dates, development security, reimbursement and termination payments should reflect the commitments made before each condition is satisfied. The review should follow the effect of a water interruption, permit restriction or other non-power condition through the energy agreements: which obligations continue, which remedies are available, and who bears costs already incurred when the customer cannot take power?

XI. The Seams Between the Contracts

Review each disruption across the wholesale and retail power contracts, toll, gas supply, pipeline agreement, hedge and guarantee: who must still perform, who can recover costs, and from whom?

The COD Seam

A commercial operation date (COD) extension clause is a financing decision, and the documents should say who pays carrying costs, committed capacity charges and bridge-power costs during the delay.

The Transportation Seam

If the customer is excused from taking power while the pipeline continues billing demand charges, the documents must allocate unused-capacity costs, release or remarketing rights, proceeds, and any obligation to preserve capacity for a restart.

The Dispatch Seam

When dispatch changes after fuel has been nominated, the toll, QSE and gas agreements must identify who relays the instruction, revises nominations, and bears imbalance and start-up costs.

The Hedge Seam

A load reduction does not reduce the hedge volume: identify the project's hedges, authority to close them, and allocation of gains and unwind losses, and test customer remedies and liability caps against those costs.

The Curtailment Seam

Mandatory curtailment and voluntary price responses need separate rules allocating program payments and avoided purchases alongside fuel, hedge and plant costs, without counting the same benefit twice.

The Change-in-Law Seam

A change in law resulting in a change in costs may be a permitted pass-through under the generator's contract but may stop at the customer's fixed-price promise: trace the change in law and any potential impacts, including affected charges, through both agreements and specify thresholds, notice, mitigation and adjustment rights.

The Credit Seam

Match each payment and termination obligation to the entity that owes it and the security that supports it, including commitments made before a guarantee begins and exposures surviving affiliate transfers, ownership changes or security reductions.

For each material event, a cross-contract schedule should identify the decision-maker, notice path, surviving obligations, cash flows, mitigation rights and final loss bearer.



Paul Turner
Turner Commodities Law, PLLC
713-899-5027
paul.turner@turnercommoditieslaw.com

James Guy
Deacon Law Group, PLLC
512-576-2435
jamesguy@deaconlawgroup.com



XII. A Practical High-Level Approval Checklist

This high-level checklist identifies principal approval issues; it is not exhaustive diligence. Each answer should identify the responsible party, relevant date, amount or formula, and controlling document.

Load

- Which MW and ramp dates are binding, and which remain forecasts or options? What evidence must the customer provide before the supplier commits capital?
- If the load arrives late, reaches only part of the forecast, or moves to another site, what minimum payments and exit rights apply?

Project

- Who controls the site and each required interconnection, construction, water, and operating milestone? Which conditions must be met before costs become irreversible?
- What may extend each commercial operation date, for how long, and who funds equipment, capacity, and bridge supply during the extension?

Gas

- Who buys fuel, nominates it, and pays for start fuel, imbalance, minimum volumes, and replacement gas when dispatch changes?
- Does the gas supply term and delivery point fit the plant's schedule and the customer's load ramp?

Transportation

- Who pays pipeline reservation charges before operation and during an outage or delayed load? Is that obligation covered by customer security?
- Who may release unused capacity, who receives the proceeds, and must the capacity remain available for a restart?

Generation

- Who may dispatch the plant, within what permit and equipment limits, and who communicates the instruction to the generator's QSE and gas shipper?
- If the plant is late, derated, or down, who obtains replacement power, who communicates that change to the load's QSE, and who bears the resulting price and congestion exposure?

ERCOT

- Who is the retail service provider and who is the QSE for the load and each resource? Who supplies forecasts and telemetry, checks settlements, and corrects errors?
- Who can commit or curtail load, receive any market payment, and pay for a failed response? Are program commitments compatible?

Hedging

- Which gas, power, basis, congestion, shape, and volume exposures remain open, and for how many years can the intended positions actually be placed?
- What liquidity is needed for margin calls, and who bears hedge gains or unwind losses if the project changes or ends?



Paul Turner
Turner Commodities Law, PLLC
713-899-5027
paul.turner@turnercommoditieslaw.com

James Guy
Deacon Law Group, PLLC
512-576-2435
jamesguy@deaconlawgroup.com



Credit

- Which entity owes each payment and termination amount, and when does its guarantee or other security become effective?
- Do liability caps and security cover the fixed investments and contingent costs already incurred before the customer takes power?

Regulatory

- Which PUCT, ERCOT, utility, water, and air approvals remain, and who bears a denial, delay, or condition that changes the planned operation?
- If a new rule or transmission charge changes cost, which agreement permits relief and can that cost reach the customer?

Failure case

- If the project is late, smaller than planned, curtailed, or never completed, what obligations continue under every agreement, and who ultimately pays for them? Have parties sought to secure appropriate rights in a future bankruptcy proceeding?

Conclusion

Hyperscale data centers offer energy companies large, potentially durable new demand. Serving that demand can require generation, gas supply, transportation, transmission, and trading commitments well before the campus reaches its expected load. The opportunity is substantial, and so are the costs that can remain if construction, energization, or the load ramp does not go as planned.

The commercial review should follow those commitments through the entire project. It should distinguish requested capacity from binding demand, test the dates and remedies in each agreement against the others, and identify who controls dispatch, curtailment, and market participation. It should also show where fuel, replacement power, hedges, and credit support land when the project changes.

The question for approval is:

If one part of this project does not happen when expected, what obligations remain everywhere else?

The parties should be able to answer with a responsible entity, a payment rule, and an enforceable source of credit support. If those answers are missing, the expected margin does not yet describe the risk of the deal.