CLIMATE CABINET EDUCATION

Taming Data Center Turmoil: Hyperscalers Push Electricity Markets to Their Limits

CLIMATE CABINET EDUCATION

JUNE 2026

Table of Contents

Executive Summary 3

Introduction 7

Why Legacy Models Break Under Hyperscale Data Center Growth Conditions 8

The Core Problem: Who Takes the Risk vs. Who Pays the Price? 10

What Failure Looks Like in Markets 11

Key Risk Signals to Watch For 14

What This Means for Legislative Action 17

Glossary of Key Terms 18

Executive Summary

Over the past century, states have built and finely tuned a regulatory apparatus to manage the power sector. These frameworks were designed for a world of gradual, predictable load growth. Today, hyperscale data centers are breaking the system. They are upending every assumption those frameworks depend on, faster than regulators can respond. More than 300 bills were introduced across 30 states in the first six weeks of 2026, underscoring the recognition that our guidelines are being overwhelmed, as well as the struggle to keep pace with hyperscale demand.

A system best equipped to manage tidal flows is facing a tsunami.

SCALE

This unprecedented situation is occurring because hyperscalers are causing six conditions to converge, creating challenges no regulatory framework was designed to manage simultaneously:

- Facilities large enough that a single exit can fundamentally alter project financials and who pays for them
- 5–10 year contracts for infrastructure built to last 30 years, creating years of uncertainty after contracts expire
- Dedicated assets with little or no alternative use
- Hyperscale developer-supplied electricity demand forecasts, shaped by a financial interest in approval and, thus, far from impartial
- Pressure to build new natural gas systems, locking in decades of climate pollution
- Globally mobile hyperscalers who can exit when conditions change, leaving consumers on the hook for stranded assets

Electricity market design affects how decisions are made, but not the ultimate financial risk to consumers and the system. Regulated utilities, competitive wholesale markets, and fully deregulated systems have each responded differently to hyperscale data center demand. Yet, they have all ended up with the same underlying outcome: concentrated infrastructure risk, opaque cost allocation, and financial exposure distributed across customers who had no role in driving the investment.

WHAT IS AT STAKE

The stakes in the hyperscale energy debate are substantial and far-reaching. They influence affordability, environmental outcomes, and the reliability of essential services. And they fall on specific people in specific places.

FOR RATEPAYERS

Households, small businesses, hospitals, and schools are effectively on the other side of every hyperscale energy deal. When a data center's demand forecast is too high, or infrastructure built for one facility is underused after it leaves, ratepayers cover the costs. Reliability interventions needed to manage the concentrated electricity demand hyperscale data centers require also flow through tariffs and surcharges that are hard to see on individual bills. Across the market structures examined in this report, ratepayers ended up paying infrastructure costs whether or not the forecast electricity demand arrived.

A closing window for legislative action. The next 12–24 months are a critical policy window. Decisions made during this period will determine how the costs and risks of the AI infrastructure buildout are allocated, and whether consumers are on the hook to pay for infrastructure that doesn't benefit them at all. Once infrastructure commitments are in place, the chance to realign who takes the risk and who pays the price is largely gone.

FOR THE CLIMATE

Hyperscale growth is pressuring utilities to build new natural gas plants. Each plant built for data center demand locks in 20–30 years of carbon pollution that future policy cannot undo. States that have cleaned up their electric grids now risk conflicts with their own climate goals because existing rules do not balance near-term reliability needs against long-term decarbonization. The window to avoid this outcome is narrow and shrinking with each new fossil fuel agreement.

FOR HOST COMMUNITIES

Data centers and the fossil-fueled plants that serve them are often sited near lower-income communities and communities of color that already face higher pollution burdens. Once facilities are permitted and operating, residents have limited tools to challenge their impacts. Laws that strip local zoning authority or restrict nuisance claims remove some of the last protections these communities have. Environmental justice is not a side issue; it follows from the fact that those who carry the financial and regulatory burden are not the ones creating the risk.

FOR GRID RELIABILITY

Hyperscale loads are very large, highly concentrated, and arrive faster than planning cycles can manage. A single campus can add 15–30% of local peak demand in a short period. If that customer later cuts usage, relocates, or shuts down, the resulting cost-recovery gap can destabilize the financial model for the entire service territory. The concentration that makes these projects attractive to developers is the same feature that turns them into systemic risks when conditions change.

FOR STATE AND LOCAL GOVERNMENTS

Governments face growing fiscal risks from hyperscale data center development that legacy frameworks are not designed to anticipate. When tenants leave or fail to materialize, stranded assets, such as transmission infrastructure, remain on utility balance sheets, shifting costs to taxpayers and ratepayers without a clear means of recovery. Tax incentives for data centers can also backfire: if facilities underperform or shut down, state and local governments are left with unrecoverable revenue shortfalls. At the same time, public infrastructure investments in roads, water systems, and grid connections create long-term obligations that persist even if facilities exit and the tax base shrinks. In effect, developers gain favorable siting and the option to walk away, while governments are left with costs that are often unanticipated and difficult to recover.

EIGHT RISKS TO WATCH FOR

We identify eight sources of risks that policymakers should consider when assessing potential data center development. These risks may primarily affect (1) the reliability of the grid and costs to ratepayers, (2) pollution harming the global climate and local communities, or (3) they may present risks overlapping both categories. In each case, policymakers can assess whether and to what extent a particular proposal presents these risks by focusing on core questions relevant to each source of risk. By identifying the primary sources of risks, policymakers can better design appropriate policy responses to mitigate and manage them.

RISK	AT RISK	CORE QUESTIONS FOR POLICYMAKERS
Single-customer concentration	Grid & Ratepayer	Does this facility exceed 15–20% of local peak demand?
Timeline compression	Grid & Ratepayer	Months
Dedicated infrastructure	Grid & Ratepayer	Local clusters with high grid impact
Forecast opacity	Grid & Ratepayer	Globally mobile
Gas dependency	Overlapping	Does this load require new fossil capacity with 30-year consequences?
Unconditional subsidies	Overlapping	Are incentives recoverable if commitments fail?
Local authority preemption	Climate & Community	Does this bill remove local permitting discretion?
Siting near vulnerable communities	Climate & Community	Has a cumulative impact assessment been conducted?

Solutions depend on which risks are present. Protecting consumers and the climate will require different responses based on how many and which of the risks above are triggered:

- **No risks triggered:** Manage through existing Public Utility Commission (PUC) processes
- **Several risks triggered:** Active legislative intervention is required
- **All four Grid & Ratepayer risks are present:** Immediate action is necessary, regardless of market structure
- **Either Overlapping risk is present:** The response must address both ratepayer and pollution dimensions, or significant financial and community risk will remain unmanaged.

THREE THINGS EVERY POLICYMAKER NEEDS TO KNOW

1. **Hyperscalers are causing structural problems for our regulatory frameworks.** Frameworks are not failing because of inadequate regulatory decisions; rather, the system is not working because its mandates were written for conditions that no longer exist. State and local lawmakers have the authority and an urgent need to update regulatory structures through new legislation, rather than narrow or surgical fixes to specific regulatory mistakes.
2. **These problems reinforce each other.** Concentration, timeline compression, dedicated infrastructure, and forecast opacity are not four separate issues. Together, they form one structural problem with four symptoms. Addressing only parts of the problem leaves the core misalignment between risk and responsibility in place.
3. **The critical moment is the point at which commitment to a hyperscale project is made.** Once a siting and permitting plan is approved, an interconnection agreement is signed, or an incentive package is granted, the financial motive to double-check the forecast is essentially gone, and risk is baked in for consumers and the climate.

The policy challenge is not load growth itself – it is the misalignment of risk and responsibility. Every hyperscaler proposal raises the same three questions:

- Who pays?
- Who benefits?
- Who bears the consequences if projections prove wrong?

If policymakers take one thing from this report, it is this: The proposals on their desk this session are not anomalies. They are the leading edge of a structural shift in electricity demand that existing regulatory frameworks were not designed to manage. The question is not whether their state will face these risks — it is whether the legislature acts before the infrastructure commitments that render them irreversible have already been made.

Introduction

The United States is entering a period of electricity demand growth unlike anything seen in decades. Unlike the gradual increases driven by traditional industrial expansion, this surge is fueled by hyperscale data centers—facilities that can consume as much power as a small city. These centers are being built in months rather than years, outpacing the regulatory frameworks designed to manage grid investment.

What makes this moment fundamentally different from past periods of industrial load growth is not just its scale, but the convergence of six conditions:

- Facilities are now large enough that the loss of a single customer can destabilize a utility's cost recovery
- Contract terms typically span 5–10 years, while the infrastructure built to serve them is financed over 30 years
- Assets are often dedicated to a single tenant, with limited alternative use if that tenant exits
- Developers frequently produce demand forecasts with a direct financial interest in project approval
- Reliability concerns have prompted renewed pressure to build natural gas capacity, reversing years of deliberate diversification away from fossil fuel-price risk
- Hyperscale operators are globally mobile. A facility can leave, abandoning infrastructure financed over decades, leaving consumers on the hook

Existing regulatory frameworks were designed to handle any one of these conditions on its own. They were not built to manage all six at once.

In just the first six weeks of 2026, more than 300 data center bills were filed across 30 states. The volume and geographic spread of proposed legislation shows that no single market design is to blame. Regulated utilities, competitive wholesale markets, and fully deregulated systems have each reacted differently to hyperscale demand, but all have converged on the same underlying outcome:

- Infrastructure risk is concentrated in a small number of facilities
- Costs are allocated in opaque ways to the customers who ultimately bear them
- Financial exposure is spread across ratepayers who played no role in driving the investment

The root cause is not that regulators have made poor decisions within their existing authority. Rather, this authority was designed for conditions that no longer exist.

This guide is written for state policymakers and staff who are navigating this mismatch in real time. It is designed to help lawmakers ask better questions, recognize when a proposal triggers risks existing law cannot manage, and understand what those risks require. Having identified these risks and needed interventions here, the next installment in this series will propose an actionable legislative framework for state lawmakers seeking to manage these risks responsibly.

Why Legacy Models Break Under Hyperscale Data Center Growth Conditions

Rapid data center expansion is revealing fundamental limits in regulatory frameworks that were built for slower, more predictable load growth. These frameworks assume that:

- Demand will rise gradually over time
- Infrastructure costs can be recovered reliably over decades
- System investments will broadly benefit many customers

Hyperscalers break each of these assumptions at once.

When that happens, the current system has predictable patterns of failure in which ratepayers bear costs whether or not the projected electricity demand materializes. In *Taming Data Center Turmoil: The Rise of the Hyperscalers*, this outcome appears in the first row, third column of the scenario matrix: cross-subsidized growth when facilities operate. Costs are spread across the full rate base, and stranded assets (partial or full) emerge when facilities underperform or exit.

HOW LEGACY MODELS FAIL

There are four main ways that legacy frameworks break down under hyperscale demand.

TIMELINE COMPRESSION AND LOAD CONCENTRATION

Large-scale loads now arrive faster than planning cycles can adapt, concentrated in specific locations and backed by developers with different risk tolerances and time horizons than traditional utility customers. A planning cycle designed for annual percentage-point growth in electricity system demand cannot adequately regulate a single facility that adds 15–20% of peak load in a single interconnection application.

PLANNING UNDER UNCERTAINTY AND COST ALLOCATION GAPS

Infrastructure decisions are increasingly made under uncertain conditions, with costs committed before long-term usage patterns are fully understood. Existing cost-allocation mechanisms were not built to distinguish between system-wide benefits and highly localized, customer-driven investments. The result is high costs spread across ratepayers who neither caused nor benefited from the new electricity demand. This is the mechanism that drives the cross-subsidized growth outcome in our *Rise of the Hyperscalers* matrix: facilities operate as expected, but the infrastructure costs associated with their load are reflected in the general rate base rather than a customer-specific cost responsibility agreement.

MISALIGNED FINANCIAL INCENTIVES

Legacy models often guarantee cost recovery once assets are approved, regardless of whether projected loads materialize. This shifts risk away from developers and onto customers, weakening disciplined decision-making around siting, efficiency, and long-term system value. **It is also the mechanism that drives the stranded-assets outcome: once infrastructure is approved and built, the financial incentive to ensure that the underlying demand forecast is accurate is no longer present.**

UNACCOUNTED SYSTEM-LEVEL RISKS

Legacy frameworks are not equipped to account for the compound risks that emerge at the intersection of reliability, cost, and environmental impact. Decisions that appear justified under a narrow economic lens may, at the same time, lock in higher long-term expenses, increase price volatility, and harm nearby communities. Gas or diesel plants built to serve a single data center are not just a question of whether the utility recovers its costs. It also locks in 30 years of carbon pollution and burdens the host community with air pollution. When frameworks treat these issues as separate questions, they overlook how these risks reinforce one another and cannot be managed in isolation.

Legacy models do not just fall short in this environment; they consistently fail to protect ratepayers. **Without updated frameworks that assign risk more precisely, require stronger load commitments, and account for system-wide impacts, data center load growth will continue to outpace the protections those models were meant to provide.**

The Core Problem: Who Takes the Risk vs. Who Pays the Price?

The core problem

The policy challenge is not load growth itself. It is the misalignment of risk and responsibility. Every hyperscaler proposal raises the same three questions:

- Who pays?
- Who benefits?
- Who bears the consequences if projections prove wrong?

Data center proposals carry two distinct but overlapping categories of risk:

1. **Grid-reliability and ratepayer risks:** This risk arises when proposals strain planning frameworks, misalign financial incentives, or shift infrastructure costs onto customers who did not drive the investment.
2. **Climate and community risks:** These risks arise when proposals lock in fossil fuel dependency, bypass local protections, or impose pollution and noise burdens on nearby residents. These impacts typically fall on lower-income communities with limited recourse.

These are not parallel concerns. Many of the most important risks sit at their intersection. For example:

- A gas dependency risk is not just about cost recovery. It is also a 20- to 30-year commitment to fossil fuel-cost volatility, carbon emissions, and air-quality impacts on host communities.
- Unconditional subsidies are not just a fiscal problem. They remove the financial incentives that would otherwise push developers toward better siting, cleaner supply, and more reliable load commitments.

Policies that address only one will leave significant exposure unmanaged.

What Failure Looks Like in Markets

The four patterns of failure described above (timeline compression and load concentration, planning under uncertainty and cost allocation gaps, misaligned financial incentives, and unaccounted system-level risks) are not theoretical. They have already produced real-world outcomes across all three major U.S. electricity market structures. The cases below show how three different types of markets have responded to hyperscale demand and how, despite very different regulatory designs, all three have ended up in the same place. Each case is linked to the scenario matrix from *The Rise of the Hyperscalers* (Table 1 below) to make that connection clear.

TABLE 1. HYPERSCALE LOAD SCENARIOS

HYPERSCALE LOAD OUTCOME	HYPERSCALER BEARS COSTS	SHARED COSTS	RATEPAYERS BEAR COSTS
Full Materialization	**Managed Growth –** Demand materializes; customer covers infrastructure costs entirely	**Cross-Subsidized Growth –** Facilities operate as expected, but some costs are absorbed indirectly by utilities/ ratepayers via rates	**Direct Subsidy –** Demand realized, but ratepayers or policy programs provide explicit financial support to cover infrastructure
Partial Materialization	**Private Stranding –** Facilities underutilized; customer absorbs costs	**Partial Subsidy Growth –** Load grows more slowly than expected; costs shared between customer and ratepayers indirectly	**Stranded Assets (Partial) –** Facilities underutilized; ratepayers bear most costs, possibly via direct support or stranded investments
Does Not Materialize	**Total Private Loss –** Customer pays for unused infrastructure	**Shared Loss –** Both customer and ratepayers absorb costs indirectly	**Stranded Assets (Full) –** Facilities absent; ratepayers or programs bear full cost of unused infrastructure

REGULATED MARKET — GEORGIA

Rapid data center growth led the Georgia Public Service Commission in 2025 to approve a major expansion of generation and transmission infrastructure, centered primarily on new natural gas capacity. Utilities pre-built assets based on demand forecasts that developers, not utilities, regulators, or independent third parties, mostly supplied. This approval creates long-lived infrastructure effectively dedicated to a small number of customers. Independent analysts warned that if projected demand fails to materialize, ordinary customers could bear high fixed costs with no practical recourse.

What this reveals about regulatory limits: Regulated systems give developers certainty but concentrate forecasting risk within the utility and spread costs across all customers. The Georgia PSC's process functioned as intended, yet still created ratepayer exposure that commissioners' interpretation of existing regulatory authority could not prevent. The lesson is that the commission's mandate did not cover the specific risks it was being asked to manage.

Scenario matrix position: Georgia is tracking toward cross-subsidized growth if projected demand materializes. Facilities would operate as expected, but infrastructure costs would be absorbed indirectly by ratepayers through rate hikes. If demand fails to materialize, the trajectory shifts toward stranded assets (in whole or in part), with ratepayers bearing most of the costs through the general rate base. No mechanism currently exists to redirect that exposure back to the hyperscale developers whose forecasts drove the investment.

COMPETITIVE WHOLESALE — PJM

A single proposed data center campus in Western Pennsylvania was projected to consume approximately 30% of local peak load (roughly 100 MW), effectively adding the equivalent of a small city such as Burlington, Vermont, to the grid almost overnight. This contributed to capacity prices hitting the market cap and prompted PJM to consider accelerated transmission projects and revised capacity rules. Though framed as technical adjustments, these interventions effectively re-centralized planning decisions that markets were intended to resolve through price signals alone: auction prices were capped, plants slated for retirement were kept online, and participation by specific resource classes was mandated. The prospective transmission projects were estimated at several hundred billion dollars, with costs distributed across every load-serving entity in a multi-state market, effectively creating an invisible line item for any individual ratepayer or legislature.

What this reveals about regulatory limits: Competition normally spreads risk and keeps costs in check. But when a single massive customer dominates demand, it triggers system responses that are difficult to track. When a data center campus strains the regional grid, PJM may need to revise its capacity rules, and the resulting costs do not fall on the data center. They are spread across all ratepayers in the region, including people in states that had no say in where the facility was built.

At every step, the system functioned as designed. The data center responded to price signals. The grid operator maintained reliability. Costs were allocated through established formulas. Nothing failed, yet no single regulator could see the costs coming, trace them to their source, or stop them.

The reason is structural. Grid operators respond to system-wide stress, not individual customers. When they adjust capacity rules, they are reacting to aggregate conditions.

Scenario matrix position: This case led to a shared-cost outcome. The facilities kept running, but the costs of market interventions were spread across many states rather than borne by the customer who created the demand. The costs appeared as capacity price events and revised tariffs rather than transparent rate increases. That is an important warning sign: in competitive markets, ratepayers can end up paying substantial hyperscale-related costs that never appear in a form that regulators or policymakers can easily see, trace back to specific projects, or challenge.

DEREGULATED, ENERGY-ONLY — ERCOT

Hyperscale growth in Texas coincided with extreme weather, tight reserve margins, and acute political sensitivity to power outages. The system responded with emergency reliability programs and increased natural gas dispatch, with costs appearing indirectly through price volatility and supplemental payments to maintain grid reliability, financed by surcharges on customer bills rather than transparent rate adjustments. In the periods following major weather events, these supplemental reliability payments totaled billions of dollars, none of which appeared as a named line item attributable to hyperscale demand on any individual customer's bill.

What this reveals about regulatory limits: Fully deregulated markets shift risk onto consumers and, ultimately, onto policymakers when reliability fails publicly. Hyperscale demand is adding concentrated, predictable baseload to a system designed for distributed, variable consumption — the opposite of what a tight-margin grid needs. When those conditions collide with an extreme weather event, the resulting costs materialize immediately. Yet under existing frameworks, it is effectively impossible to trace those costs back to the specific load decisions made years earlier.

Scenario matrix position: Texas is the most complex case to place in *The Rise of the Hyperscalers* matrix because costs appear through price volatility and customer surcharges rather than transparent rate adjustments — a structural feature of energy-only markets that the matrix does not fully capture. This is itself a significant diagnostic finding: the scenario matrix assumes that costs, when they materialize, will be traceable. In deregulated markets, that assumption does not hold. Requirements that cannot be traced cannot be enforced.

WHAT THESE CASES MEAN TOGETHER

Market design affects how decisions are made, but not the ultimate financial risk. Regulated utilities, competitive wholesale markets, and fully deregulated systems have each responded differently to hyperscale data center demand. **Yet, they have all ended up with the same underlying outcome: concentrated infrastructure risk, opaque cost allocation, and financial exposure distributed across customers who had no role in driving the investment.**

This is not a failure of any individual market design. It shows that the core assumptions all three models share (gradual load growth, shared infrastructure, auditable forecasts, matched asset and contract lifetimes) no longer reflect the conditions under which those markets now operate. **When the assumptions underlying a model no longer hold, the model stops producing reliable results, even if everyone follows the rules. That mismatch is the structural problem that legislative action has to solve.**

The eight risks in Table 2 are the main places where this structural gap is likely to show up in a real proposal. They signal whether a proposal is something the current regulatory framework can handle.

Key Risk Signals to Watch For

Policymakers reviewing active proposals are most likely to encounter eight specific risks, and these are precisely the ones that the existing regulatory environment is least equipped to manage.

The first four — **concentration, timeline compression, dedicated infrastructure, and forecast opacity** — are invisible under the legacy planning model. They only emerge when examining a single customer's relationship to the system as a whole, something old frameworks were never designed to do.

The next two — **gas dependency and unconditional subsidy** — are frequently miscategorized. Because each straddles more than one policy domain, they tend to draw only partial legislative responses, leaving important dimensions unaddressed.

The final two — **local authority preemption and siting near vulnerable communities** — represent the protections most accessible to affected residents. Once removed, they are among the hardest to restore.

Together, these eight risks mark the boundary between proposals that the current regulatory framework can manage and those requiring the updated approach described in Part Three of the *Taming Data Center Turmoil* series.

How to use this framework:

- Proposals that trigger none of these risks may be handled through existing PUC processes.
- Proposals that trigger several, especially any dual-risk combination, require active legislative intervention.

The table below presents all eight risks at a glance. The Risk Type column notes whether each risk is mainly a grid and ratepayer issue, a climate and community issue, or both. The Red Flag Link column points to the matching warning signal from *The Rise of the Hyperscalers*. The Questions to Ask column is meant for practical use in hearings, negotiations, and bill review.

TABLE 2. EIGHT KEY RISKS IN HYPERSCALE ENERGY PROPOSALS

RISK	AT RISK	WHAT TRIGGERS IT	RED FLAG LINK	QUESTIONS TO ASK
Single-customer concentration	Grid and ratepayer	A proposed load represents more than 15–20% of current peak demand in its service territory[1]	Customer Concentration: transmission or generation upgrades justified by a single customer	What share of current peak demand does this facility represent? What happens to fixed cost recovery if this customer exits?

[1] Note: the 15–20% threshold is the benchmark illustrated by the Western Pennsylvania case study, in which a single campus at approximately 30% of local peak triggered market-cap events and re-centralized planning.

RISK	AT RISK	WHAT TRIGGERS IT	RED FLAG LINK	QUESTIONS TO ASK
Timeline compression	Grid and ratepayer	A utility is committing to infrastructure with a 20–30 year useful life to serve a customer whose contract runs 5–10 years or less	Infrastructure Timing: built before load commitments are confirmed	What is the term of the interconnection agreement? Who bears the cost if the customer exits at the end of the contract?
Dedicated infrastructure	Grid and ratepayer	New substations, transmission lines, or generation capacity would have limited alternative use if the customer reduces load or relocates	Customer Concentration: the hyperscale customer should fully bear the cost of triggered infrastructure	Could this infrastructure serve other customers if the data center exits? Has the utility modeled alternative-use scenarios?
Forecast opacity	Grid and ratepayer	Demand projections rely heavily on a single developer's internal load estimates rather than independent analysis	System-Wide Planning: clustering impacts assessed project-by-project rather than system-wide	Who produced the demand forecast? Was it independently verified? What are the downside scenarios, and who modeled them?
Gas dependency and fossil fuel facilitation	Both	Meeting the new load requires significant new natural gas capacity; OR a bill sites, permits, subsidizes, or expedites new gas generation, peakers, or cogeneration primarily to serve data center demand	Long-Term Implications: near-term reliability decisions locking in expensive or inflexible infrastructure	What generation mix is proposed? Does the bill create pathways for new fossil infrastructure? How does this affect the utility's decarbonization commitments and host community air quality?
Unconditional subsidies	Both	Tax exemptions, abatements, or incentives with no clean energy, ratepayer protection, EJ, or load commitment conditions attached	Cost Recovery: costs recovered from ratepayers before demand materializes	What conditions, if any, are attached to any incentive? Is there a clawback mechanism? What public benefit is the subsidy purchasing?

RISK	AT RISK	WHAT TRIGGERS IT	RED FLAG LINK	QUESTIONS TO ASK
Local authority preemption	Climate and community	Provisions removing local zoning discretion, restricting noise or environmental standards, limiting local permitting authority, or preempting nuisance claims by affected communities	Contracts and Processes: special contracts that bypass standard planning processes	Does the bill restrict what local governments can require of or from data center developers? Does it limit community legal remedies?
Siting near vulnerable communities	Climate and community	A proposed facility is sited near low-income communities or communities of color without a cumulative impact assessment or community benefits process	Contracts and Processes: absence of transparent public review	Has a cumulative pollution burden assessment been conducted? Is there a binding community benefits agreement? Are diesel backup generator restrictions included?

The first four risks operate as an interconnected system. When one large customer on a short contract needs its own dedicated infrastructure based on its own demand forecast, it breaks all the core assumptions of the legacy planning model at once:

- That demand grows gradually and predictably over time
- That investments provide benefits across multiple customers
- That forecasts can be independently audited and verified
- That customer commitments align with the lifespan of the infrastructure built to serve them

The two dual risks—gas dependency and unconditional subsidies—are especially important to identify early because they are often misclassified. A policymaker focused on ratepayer protection, for example, may see gas dependency mainly as a cost-recovery issue and miss the long-term impact of 30 years of climate pollution. Focusing only on the financial side may protect ratepayers in the short term but leaves long-term climate and infrastructure risks unmanaged. The reverse is also true: focusing only on infrastructure and climate concerns without ensuring fiscal accountability can shift costs to consumers in unpredictable ways. **Effective responses must do both: limit financial exposure and prevent long-term carbon pollution.**

What This Means for Legislative Action

Working through the eight risks outlined in Table 2 against an active proposal answers the first set of questions a policymaker needs to answer: **Is this proposal within the capacity of existing regulatory frameworks, or does it require active intervention?** If the answer is intervention, the next question is what kind, in what sequence, and calibrated to which market structure?

Those questions are answered in the next installment in this series, which provides the nine-element legislative framework and its tier structure, and in the subsequent installment, which includes the pathway selector, case studies, and implementation sequence by market type. The diagnostic work in this document is the necessary foundation for using those tools well. A policymaker who has worked through Table 2 against its active proposals will know which elements are most urgent, which risks are active, and which industry arguments are likely to be encountered.

Three things this diagnostic should have made clear before moving on to the prescriptive guides:

1. **The regulatory gap is structural, not incidental.** The patterns of failure in this document are not the result of bad regulatory decisions. They stem from a regulatory environment that was built for conditions that no longer exist. **Legislative action is needed not because current agencies have performed poorly, but because their legal mandates do not cover the risks created by hyperscale development.**
2. **The risks interact.** A proposal that triggers single-customer concentration, timeline compression, dedicated infrastructure, and forecast opacity at the same time is not four separate problems. It is one structural problem with four symptoms. Lawmakers need to fix the underlying mismatch between who takes the risk and who is responsible, rather than treating each symptom in isolation.
3. **The window for prevention is before the commitment, not after.** The most important decisions — Integrated Resource Plan (IRP) approvals, siting and permitting decisions, interconnection agreements, and incentive packages — are made before the public can see how much risk ratepayers will face. Once the infrastructure is approved and built, there is no longer a strong financial incentive to ensure the original demand forecast was accurate. The legislative framework in the next installment in this series is designed to act at this point of commitment, before those long-term risks are locked in.

To lawmakers, if you take one thing from this diagnostic: The proposals on your desk this session are not anomalies. They are the leading edge of a structural shift in electricity demand that existing regulatory frameworks were not designed to manage. The question is not whether your state will face these risks — it is whether your legislature acts before the infrastructure commitments that render them irreversible have already been made.

Glossary of Key Terms

Carbon Lock-In: The long-term commitment to carbon-emitting infrastructure that results from building fossil fuel generation capacity to serve new demand. Because power plants operate for 20–30 years, decisions made today constrain decarbonization options for decades.

Clawback Mechanism: A contractual or statutory provision that requires a subsidy recipient to return some or all of the benefit received if specific conditions—such as job creation targets, load commitments, or clean energy requirements—are not met.

Clustering: The concentration of multiple data center facilities within the same geographic area or transmission zone. While individual proposals may appear manageable in isolation, clustering can produce cumulative grid impacts that project-by-project analysis fails to capture.

Community Benefits Agreement (CBA): A legally binding contract between a developer and a community organization that specifies commitments the developer must fulfill in exchange for community support—such as local hiring, environmental mitigation, or funding for public services.

Cost Socialization: The spreading of infrastructure costs across a broad base of ratepayers, including customers who did not drive and do not benefit from the investment. Cost socialization is a central concern when dedicated infrastructure is approved, but its costs are treated as system-wide.

Cumulative Impact Assessment: An evaluation of the combined pollution burden—from air quality, noise, water, and other sources—already experienced by a community, used to assess whether a new facility would compound existing harms.

Dedicated Infrastructure: Transmission lines, substations, or generation capacity built primarily or exclusively to serve a single customer. If that customer exits, dedicated infrastructure may have limited alternative uses, creating stranded-asset risk.

Demand Forecast: A projection of the electricity load a proposed facility is expected to generate. Forecasts produced by developers with a financial interest in project approval may differ significantly from independent analyses, particularly under downside scenarios.

Environmental Justice (EJ): A policy framework that addresses the disproportionate exposure of low-income communities and communities of color to environmental harms, including air pollution, noise, and industrial development. EJ concerns are particularly relevant when data centers with large diesel backup generator fleets are sited near residential areas.

Hyperscale Data Center: A facility typically exceeding 100 megawatts of IT load capacity, operated by or on behalf of major technology companies for cloud computing, artificial intelligence, and large-scale data processing. These facilities are distinguished from smaller commercial data centers by their size, power consumption, and the speed at which they are developed.

Interconnection Agreement: A contract between a utility and a large electricity customer—or a generator—specifying the terms under which the customer connects to the grid. The term length of these agreements is a key variable in assessing stranded asset risk.

Load Commitment: A binding obligation by a large electricity customer to consume a specified minimum level of power over a defined period. Strong load commitments shift financial risk back toward the developer and away from ratepayers.

Peaker Plant: A power generation unit that operates only during periods of peak electricity demand. Peakers are typically natural gas-fired and carry higher per-unit costs than baseload generation. New peaker construction to serve data center demand raises both cost and carbon lock-in concerns.

Public Utilities Commission (PUC): The state regulatory body responsible for overseeing investor-owned utilities, approving rate structures, and authorizing major infrastructure investments. PUCs operate under authority granted by state legislatures and their processes are a primary venue for addressing hyperscale energy risks.

Ratepayer: Any customer who pays electricity bills under rates set by a utility and approved by a PUC. Ratepayers—residential, commercial, and industrial—ultimately bear the costs of infrastructure investments socialized across the rate base.

Service Territory: The geographic area within which a utility has the exclusive right and obligation to provide electricity service. The share of peak demand a single customer represents within a service territory is a key indicator of concentration risk.

Single-Customer Concentration: A condition in which one customer accounts for a disproportionate share—generally 15–20% or more—of a utility's peak demand. High concentration creates vulnerability: if that customer exits or reduces their load, fixed-cost recovery may be destabilized for all remaining ratepayers.

Stranded Asset: Infrastructure that has been financed and built but that earns less revenue than projected—often because the customer it was built to serve has exited, reduced consumption, or failed to materialize as expected. The costs of stranded assets are typically recovered from ratepayers.

Tax Abatement: A reduction or elimination of property or other taxes offered as an incentive to attract development. When offered without conditions or clawback provisions, abatements reduce public revenue while providing no guarantee of the economic activity they are meant to secure.

Transmission Upgrade: An expansion or reinforcement of high-voltage power lines and associated infrastructure needed to deliver electricity to a new large load. When upgrades are required primarily for a single customer, cost-allocation rules determine whether those costs are borne by the customer or spread across the broader rate base.