CLIMATE CABINET EDUCATION

Taming Data Center Turmoil: A Vision for Data Center Regulation that Protects Consumers

CLIMATE CABINET EDUCATION

JULY 2026

Table of Contents

Executive Summary 3

Introduction 6

Building a Hyperscale-Ready Framework 8

The Hyperscale Framework: 9 Policy Elements to Ensure Projects Pass 3 Key Tests to Protect Consumers, the Environment, and Public Health 11

Test 1 — Before It's Approved. *11*

Test 2 — While It Runs. *14*

Test 3 — After the Deal Is Done. *16*

Existing Facilities — A Different Legislative Problem 20

Enforcement — A Cross-Cutting Requirement 21

Glossary 27

Executive Summary

Hyperscale data centers are arriving faster, larger, and more concentrated than the institutions built to manage them. A single campus can demand a small city's worth of electricity on a timeline of months, infrastructure investments that last decades, and contracts that expire long before assets are paid off.

This pace and scale of development is pushing regulators to make decisions shrouded in uncertainty. When data center developers supply electricity demand forecasts that overstate true need, or when a facility scales back or exits, the cost of underused or unneeded infrastructure falls on the shoulders of consumers who didn't ask for or benefit from the project.

Our current system for regulating hyperscale data centers is simply unacceptable– there's a structural power imbalance favoring data center developers over communities. Data center developers are best positioned to know whether a project will deliver, and best able to absorb the losses if it does not. Yet they carry the least exposure to its failure. A developer can overstate electricity demand, trigger hundreds of millions of dollars in infrastructure commitments, generate millions of tons of pollution and massive public health impacts, and just walk away if projections fall short, despite everything that has been done to accommodate them. The utility recovers its investment through rates either way. And the families and businesses who had no seat at the table are left to deal with the consequences of decisions they never made.

None of these bad outcomes require bad faith– they are simply what existing rules produce when growth at this scale and speed meets institutions designed for a slower, more predictable era. This paper offers policymakers a vision to right these wrongs.

Every risk this paper addresses comes down to two questions: Will projected electricity demand materialize, and who will bear the costs of the infrastructure built to serve the data center?

This paper offers policymakers a framework to regulate data centers that protects consumers and the climate. Its nine policy elements are organized around three policy tests that follow a data center through its life — before it is approved, while it runs, and after the deal is done. Passing these tests in order, with all policy elements included before consumer money is at risk, is the protective response to the risks identified in the companion paper, *Taming Data Center Turmoil: Hyperscalers Push Electricity Markets to Their Limits*.

If a proposed data center doesn't clearly pass all three tests, decisionmakers risk shifting costs and other impacts onto their constituents for the benefit of extremely profitable corporations.

More than 300 bills on data center regulation were introduced across the country in just the first six weeks of 2025, proof that lawmakers see a problem and are struggling to solve it. This framework compiles best practices from these bills, and offers policy proposals for filling remaining gaps. Importantly, it *does not* ask states to choose between attracting data center investment and protecting families and businesses. States that have adopted comparable requirements continue to attract hyperscale development– **what changes is not the volume of investment but its terms.**

The goal is straightforward: ensure that whichever electricity demand scenario ultimately materializes, risks are identified early, costs are borne by those who create them, and ordinary customers are never left holding the bag for private investment decisions.

What Good Data Center Policy Looks Like

Climate Cabinet's framework defines three interdependent tiers. Policy that addresses only one or two leaves the structural risks intact.

Test 1 | Before it's approved

Has the project proven its electricity demand and put its own money behind it?

To pass this test, policy makers should hold project sponsors to standards that include:

- **Element 1** Verified Electricity Demand Forecasting
- **Element 2** Committed-Capital Requirements
- **Element 3** Cause-and-Pay Interconnection

Test 2 | While it runs

Does the facility bear the real costs and burdens it imposes?

Policymakers can ensure projects pass this test with:

- **Element 4** Real-Cost Pricing
- **Element 5** On-Call Load Flexibility
- **Element 6** Scenario-Based Planning

Test 3 | After the deal is done

Does the facility leave the public better off over the decades after its contracts expire?

Policymakers can help developers pass this test by establishing:

- **Element 7** New-and-Matched Clean Energy
- **Element 8** Community-Protective Siting Standards
- **Element 9** Performance-Tied Incentives

The policy elements that enable projects to pass these tests are designed to work together, not in isolation. Verified demand baselines make financial commitments enforceable; financial commitments give adaptive planning credible inputs; pricing signals reinforce clean energy and siting requirements without separate mandates. For each element, the paper provides a plain-language explanation, a description of what it looks like in practice, core legislative requirements, and a candid account of the tradeoffs, costs, and likely sources of support and opposition.

Two cross-cutting issues complete the framework. First, because the largest loads on many systems already operate under legacy agreements, the paper outlines three mechanisms for extending these standards to existing facilities. Second, every element depends on enforcement: independent verification, penalties calibrated to exceed the cost of non-compliance, consumer standing to intervene, and appeals processes that cannot be used to stall. **A requirement without consequences is a voluntary commitment by another name.**

CASE STUDY: VIRGINIA'S MISSED OPPORTUNITIES IN 2025 AND 2026

A case study examining Virginia, the world's leading data center hotspot, shows how this framework can expose regulatory gaps. Rather than evaluating individual bills in isolation, it can be used to assess Virginia's entire legislative portfolio. Across two sessions, Virginia produced an unusually large volume of legislation and, in 2026, enacted several measures that failed a year earlier. Yet significant policy gaps remained across every stage of the project life cycle. Applied across the 2025 and 2026 sessions, the framework's three tests yielded three findings:

- **Before approval, the stage that matters most, coverage advanced, but the portfolio fell short.** In 2025, the bills that would have verified demand, placed developer capital at risk, or assigned infrastructure costs to the customers driving them nearly all failed. In 2026, legislation establishing the principle of cost allocation for high-load customers was enacted, although subsequent amendments replaced mandatory cost-shifting provisions with more discretionary language, and independent demand verification remained optional.
- **During operation, safeguards remained only partially developed across both sessions.** The measures that passed took the form of studies, pilots, and procedural reforms rather than operational requirements. In both 2025 and 2026, Virginia did not enact real-cost pricing, impose mandatory curtailment before households, or conduct downside stress testing tied to large new loads, even as regional reliability pressures intensified.
- **Over the long term, protections expanded but stopped short.** Legislation establishing community siting requirements and water-use transparency, which did not advance in 2025, was enacted in 2026, expanding the state's oversight of data center development. Yet the state's multi-billion-dollar tax incentive remained largely disconnected from performance, and new demand was still not matched to new clean generation.

A more detailed analysis is presented in the section *Evaluating a Legislative Portfolio—Virginia's 2025 - 2026 Data Center Docket*.

Introduction

Electricity market regulations were built for a predictable world where demand grows gradually, costs recover steadily, and infrastructure investment delivers broad, shared benefits. Hyperscale data centers shatter all those assumptions at once. They arrive in months rather than years. They cluster in specific locations rather than spreading the load across the system. Their contracts expire before the infrastructure built to serve them does. And they are backed by developers who have a direct financial stake in the demand forecasts that justify that infrastructure in the first place.

The result is a systematic misalignment of risk and responsibility. The parties best positioned to know whether a project will deliver — and best able to absorb the losses if it does not — carry the least exposure to its failure. A developer can overstate demand, trigger hundreds of millions of dollars in often highly-polluting infrastructure commitments, and walk away when projections fall short. The utility recovers its investment through rates either way. And the families and businesses who had no seat at the table become the default backstop for decisions they never made. None of this requires bad faith; it is simply what existing rules produce when load at this scale and speed meets institutions designed for a slower, more predictable era.

Realigning that risk and responsibility is the purpose of this paper. This paper offers a Hyperscale-Ready Framework consisting of nine policy elements. It organizes them around three policy tests that a lawmaker can use to assess any project — or any bill — that track the stages of the infrastructure lifecycle:

Test 1 — Before it's approved.

Has the project proven its electricity demand and put its own money behind it?

To pass this test, policy makers should hold project sponsors to standards that:

- Lock in verified commitments before infrastructure is approved, with independent audits of developer demand forecasts (Element 1);
- Put in place bonds and take-or-pay contracts that put developer capital at risk ahead of consumer dollars (Element 2);
- And adopt cause-and-pay interconnection rules that assign grid upgrade costs to the customer who triggers them (Element 3).

Test 2 — While it runs.

Does the facility bear the real costs and burdens it imposes?

Policymakers can ensure projects pass this test with:

- Locational and time-varying pricing that reflects real system costs (Element 4);
- Mandatory demand response that requires hyperscale facilities to curtail before residential customers during grid emergencies (Element 5);
- And scenario-based planning that stress-tests every major investment against the possibility that demand falls short (Element 6).

Test 3 — After the deal is done.

Does the facility leave the public better off over the decades after its contracts expire?

Policymakers can help developers pass this test by establishing:

- Clean energy requirements built on additionality and time matching so hyperscale growth drives new clean supply rather than fossil-fuel lock-in (Element 7);
- Community-protective siting standards with mandatory impact studies and public participation (Element 8);
- And performance-tied incentives with clawback provisions so public subsidies flow only to facilities that deliver as promised (Element 9).

The questions are sequential but also interdependent. A verified demand forecast (Test 1) gives the pricing and planning tools of Test 2 something reliable to act on. Real financial commitments (Test 1) give scenario-based planning something firmer to work with than developer projections. And the clean energy requirements of Test 3 depend on siting rules that put projects where clean power is available.

No single element is enough on its own; the framework functions because the elements reinforce one another. Bad forecasts are caught before they trigger infrastructure spending. Financial structures that would shift private risk onto consumers are blocked. Planning cycles are rebuilt to keep pace with concentrated demand, and incentive programs that reward mere announcements are reformed to reward actual delivery. The following section presents the entire framework; each element is then examined in detail.

Building a Hyperscale-Ready Framework

The Hyperscale-Ready Framework is not a single policy—it is a coordinated set of actions designed to prevent the systemic failures that arise when large-scale load growth outpaces existing regulatory and planning structures. At their core, these risks can be distilled into two fundamental questions: **Will projected demand actually materialize, and who will bear the costs of the infrastructure built to serve it?** The answers determine whether investment decisions support reliable system growth or leave households exposed to stranded costs and underutilized infrastructure.

The first paper in this series, *Taming Data Center Turmoil: The Rise of Hyperscale Load* explains how different combinations of factors—from full electricity demand materialization, with costs borne by the hyperscaler, to complete non-materialization, with costs shifted to consumers—produce markedly different outcomes.

The risks identified in our second paper, *Taming Data Center Turmoil: Hyperscalers Push Electricity Markets to Their Limits*, define the challenge.

The nine framework elements offered here define the response to these risks. The Hyperscale-Ready Framework offered here is designed to ensure that, regardless of the scenario that emerges, risks are identified early, costs are assigned appropriately, and consumers are protected from liabilities arising from private investment decisions.

Hyperscale Load Scenarios

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

Together, they address the main ways the current policy environment allows developers to shift costs and risks to consumers, including:

- Infrastructure investments based on unverified demand forecasts;
- Financial arrangements that transfer private risk to the public;
- Planning processes that cannot keep pace with rapid load growth; and
- Incentive programs that reward project announcements without ensuring that promised investments, jobs, or other public benefits are ultimately delivered.

Our framework consists of three tests and the policy elements that lawmakers can ensure are present for a project to pass those tests. **Test 1 (before approval)** is answered by demand forecasting, committed capital, and cause-and-pay interconnection (Elements 1–3). **Test 2 (while it runs)** is answered by real-cost pricing, on-call load flexibility, and scenario-based planning (Elements 4–6). **Test 3 (after the deal)** is answered by new-and-matched clean energy, community-protective siting, and performance-tied incentives (Elements 7–9).

The framework is organized around three policy tests, each tied to a stage of the infrastructure lifecycle. Test 1 sets the conditions that must be met before any commitment is made. Test 2 governs how costs and risks are managed once development is underway. Test 3 ensures the system can adapt over time as demand and circumstances change. Each element is presented in three parts: a plain-language explanation of its purpose, a description of what it looks like in practice, and the core legislative requirements a policymaker would need to enact.

Policymakers can use the cross-reference table below to see which of the proposed policy Elements from the framework this paper offers address which risks are present.

Risk-to-Element Cross-Reference

Element	Primary risks addressed	Supporting risks addressed
Element 1 Verified Electricity Demand Forecasting	• Timeline compression • Forecast opacity	• Single-customer concentration
Element 2 Committed-Capital Requirements	• Single-customer concentration	• Dedicated infrastructure • Unconditional subsidies
Element 3 Cause-and-Pay Interconnection	• Dedicated infrastructure	• Timeline compression
Element 4 Real-Cost Pricing	—	• Gas dependency and fossil fuel facilitation
Element 5 On-Call Load Flexibility	—	—
Element 6 Scenario-Based Planning	—	• Timeline compression • Forecast opacity
Element 7 New-and-Matched Clean Energy	• Gas dependency and fossil fuel facilitation	• Siting near vulnerable communities
Element 8 Community-Protective Siting Standards	• Local authority preemption • Siting near vulnerable communities	• Single-customer concentration • Gas dependency and fossil fuel facilitation
Element 9 Performance-Tied Incentives	• Unconditional subsidies	—

The Hyperscale Framework:
9 Policy Elements to Ensure Projects Pass 3 Key Tests to Protect Consumers, the Environment, and Public Health

TEST 1 — BEFORE IT'S APPROVED.

Has the Project Proven Its Projected Electricity Demand and Put Its Own Money Behind It?

The first test focuses on the point where a developer's proposal becomes a regulatory commitment. Its elements are the most consequential in the framework. Decisions made at this stage determine whether the risks identified in *Taming Data Center Turmoil: Hyperscalers Push Electricity Markets to Their Limits* ever materialize. A full answer requires all three elements to pass the test: an independently verified forecast, developer capital at risk before consumer money, and grid upgrade costs assigned to the customer who triggers them.

ELEMENT 1:
Verified Electricity Demand Forecasting

What it targets

This element is the framework's first line of defense against the partial and non-materialization electricity demand rows in the scenario matrix laid out in *Taming Data Center Turmoil: The Rise of Hyperscale Load*. These scenarios are outcomes in which demand falls short of projections and costs have nowhere to go but to consumers). **Independently verifying developer electricity demand forecasts before any infrastructure is approved or built helps avoid the worst-case outcomes in the scenario matrix.** If the baseline shows demand projections are unrealistic, policymakers can require stricter financial commitments, price the risk correctly, or halt the investment before any consumer money is on the line. Without a verified baseline, financial commitments lack a reliable target, and planning scenarios have no credible starting point for measuring shortfalls.

What it looks like in practice

A designated independent body (housed within the Public Utility Commission, a state energy office, or established as a standalone office) audits large-load demand forecasts before any infrastructure commitment is approved. The developer's internal projections are checked against independent analyses, comparable facility data, and conservative downside scenarios before they are incorporated into an Integrated Resource Plan (IRP), capacity auction, or transmission plan.

Core legislative requirements

- **Impact studies before approval.** Before a data center can connect to the grid, the state utility regulator must complete a study of the actual cost and demand of that connection. There is a size threshold; only facilities above a certain wattage are subject to this requirement, so smaller customers are not burdened by the same process.
- **An independent auditor with clear rules.** Someone outside the developer's organization must check the numbers — a permanent office or a case-by-case appointee — against defined standards: whether the engineering works, how similar facilities have performed elsewhere, and what happens if demand comes in lower than projected.
- **Facilities must report what they actually use.** Once operational, a facility must disclose actual electricity consumption, creating a feedback loop that allows regulators to compare what was promised with what was delivered to sharpen scrutiny of future forecasts.
- **A public record of grid costs.** Every time a data center drives new infrastructure spending, that cost is recorded and published, building an evidence base of what large customers actually cost the system.
- **Deadlines that can flex without becoming indefinite.** The audit process should have a clear timeline; complex proposals may receive limited extensions that cannot be used to run out the clock on review.

ELEMENT 2:
Committed-Capital Requirements

What it targets

This element is the primary mechanism for moving outcomes from the "Consumers Bear Costs" column to the "Hyperscaler Bears Costs" column. If a facility underperforms or shuts down early, the upfront financial commitments mean recovery flows from the developer first, before any costs can reach the broader rate base. Without this element, developers could abandon underperforming projects, leaving ordinary customers to pay for stranded assets.

What it looks like in practice

Before a utility breaks ground on dedicated infrastructure, the data center developer posts a bond or letter of credit equal to a defined share of the infrastructure cost. The interconnection agreement includes a take-or-pay clause requiring the developer to pay for contracted electricity generating capacity whether or not it is consumed. If the developer exits, stranded asset costs are recovered from the bond and any remaining obligations before a single dollar flows to the general rate base.

Core legislative requirements

- **A separate rate class for large customers.** A distinct rate category ensures hyperscale bills reflect the true cost of serving their unique and outsized demand, rather than spreading those costs across the general customer base.
- **Take-or-pay contracts.** Before connecting, a large facility must commit to pay for its contracted electricity generating capacity whether or not it uses it, preventing a developer from reserving capacity, driving investment, and walking away if the project underperforms.

- **Cost responsibility agreements.** Where a connection requires dedicated infrastructure (a new substation or transmission line built solely to serve the facility), the customer must sign a formal agreement accepting financial responsibility for it.
- **Stranded asset rules.** If dedicated infrastructure goes unused, costs are first recovered from the customer who triggered the investment, in accordance with a defined order of recovery that protects ordinary customers.

ELEMENT 3:
Cause-and-Pay Interconnection

What it targets

Addresses the customer concentration and infrastructure timing red flags by ensuring that grid upgrade costs triggered by a new load are borne by the customer driving them, and that feasibility is confirmed before commitments are made.

What it looks like in practice

A developer seeking connection funds a pre-application study, completed before any interconnection agreement is signed, to assess grid impacts, required upgrades, and estimated costs. Any transmission or substation infrastructure primarily serving the new load is paid for by the developer rather than socialized across other customers. The agreement also includes minimum capacity utilization commitments and financial penalties for sustained underuse.

Core legislative requirements

- **Developers pay for their own infrastructure.** Where a new facility requires upgrades to the grid, such as a new substation, additional transmission capacity, or reinforced distribution lines, the developer funds those upgrades directly. Costs that exist solely because of a new customer's demand do not enter the rate base.
- **Studies before applications.** Before a developer can formally apply for connection, an independent study must assess what the proposed facility will actually require from the grid and what it will cost to deliver. This prevents developers from locking in agreements and triggering infrastructure commitments before the full picture of grid impact is known.
- **Minimum use commitments.** The interconnection agreement includes a binding commitment to use a specified minimum level of capacity throughout the term of the agreement. A developer who connects, drives infrastructure investment, and then operates well below projected load remains financially responsible for the capacity they reserved.

TEST 2 — WHILE IT RUNS.

Does the Facility Carry the Real Costs and Burdens It Imposes?

With the conditions of approval in place, the second question shifts from preventing bad commitments to ensuring the grid operates efficiently under the load those commitments generate — and that the facility, not everyone else, carries the real costs and burdens it creates. Passing the test requires all three associated policy elements: prices that reflect real system costs, a binding obligation to curtail electricity demand before households, and planning that is stress-tested against the load falling short.

ELEMENT 4:
Real-Cost Pricing

What it targets

This element addresses the system-wide planning gap by making the "Hyperscaler Bears Costs" column the path of least resistance. When prices reflect actual grid conditions, including where a facility is located and when it draws power, data center developers have a direct financial incentive to site projects where electricity generating capacity already exists, and to shift flexible workloads away from peak electricity demand periods. This reduces clustered infrastructure investment that skews outcomes toward stranded-asset cells and encourages buying clean energy when and where it's cheapest, reinforcing Element 7's time-matching requirements.

What it looks like in practice

A hyperscale facility in a congested electricity corridor pays substantially higher peak-hour rates than the same facility at a less constrained location with nearby available generation. A developer choosing between the two sites can see from the rate structure itself that the constrained location carries a significant ongoing cost premium, creating a strong financial incentive to site where the grid has excess capacity rather than where real estate is cheapest.

Core legislative requirements

- **Nodal or locational pricing for large loads.** Instead of a flat rate, large facilities pay prices based on the real cost of delivering power to their location. Connecting where the grid is stressed costs more than connecting where capacity is available.
- **Time-varying rates for large customers.** Large facilities pay more during peak periods when the grid is stressed and less when capacity is available; shifting non-essential work to off-peak hours reduces grid strain.
- **Transparent cost allocation for dedicated infrastructure.** Where a large customer's location or timing requires infrastructure that would not otherwise be built, the cost of that infrastructure is visible, assigned, and recovered from the customer who triggered it, rather than averaged across the rate base.
- **Regular rate review for large-load customers.** As grid conditions change (e.g., new generation comes online, demand shifts, infrastructure is built or retired), the prices large customers pay are reviewed and updated on a defined schedule to keep prices in line with real conditions.

ELEMENT 5:
On-Call Load Flexibility

What it targets

Works across the partial and non-materialization rows of the scenario matrix by converting a potential liability into a reliability asset. A facility that has underdelivered on its forecast still holds value to the grid if it can shed load on short notice during peak stress. Demand response turns that flexibility into a formal commitment, so even when demand does not materialize as projected, the facility contributes to grid stability rather than simply representing stranded investment. It is one of the few elements that improve outcomes across multiple cells of the matrix at once.

What it looks like in practice

During a grid emergency, the operator issues a curtailment signal. Enrolled hyperscale facilities must reduce load by a specified amount within 30 minutes, ahead of any residential curtailment. Facilities that exceed their voluntary floor receive compensation; those that fail their mandatory obligation face penalties. The protocol is pre-negotiated, tested annually, and verified by the grid operator.

Core legislative requirements

- **Large facilities must participate in demand response.** Any facility drawing power above a defined threshold is automatically enrolled; participation is not optional. The customers with the greatest capacity to ease grid stress are available when it matters most.
- **Clear rules for how and when curtailment works.** Each enrolled facility must meet specified requirements for how quickly it reduces load, by how much, and for how long, so the grid operator knows exactly what it can count on before an emergency arrives.
- **Hyperscale facilities are first in line for curtailment.** Under stress, large facilities reduce use before residential customers. The customers with the greatest individual impact on grid strain bear that responsibility first; residential customers face reductions only as a last resort.

ELEMENT 6:
Scenario-Based Planning

What it targets

Addresses the entire partial- and non-materialization rows by requiring utilities to model what happens before committing to infrastructure to serve demand that may never arrive. Where other elements help prevent worst-case scenarios, this one ensures utilities are ready if those situations still occur. Stress-testing major decisions against lower-than-expected demand, facility closures, and non-materialization lets utilities make smarter investments, set better contract terms, and avoid saddling consumers with costs based on overly optimistic forecasts. Without it, the potential for partial or missed demand remains unmanaged, leaving policymakers without early warnings when projections fall short.

What it looks like in practice

A utility whose territory adds 300 megawatts (MW) or more of hyperscale demand in a single year is required to file an updated IRP within 12 months, including three demand scenarios: a base case, an upside case assuming further growth, and a downside case in which two of the new customers exit within 10 years. No new generation or transmission commitment may be approved unless it is

financially viable under the downside scenario. Where a proposed investment fails that test, the PUC must make explicit findings identifying why the consumer exposure is justified before approval can proceed.

Core legislative requirements

- **Forecasts must be updated regularly and triggered by large load changes.** Utilities update demand forecasts every two to three years at minimum, and immediately whenever a new large connection exceeds a defined threshold. A forecast accurate when filed can become misleading within months if significant load arrives or departs.
- **Every forecast must include a downside scenario.** Alongside the primary projection, utilities model scenarios in which one or more large customers leave or sharply cut usage in every forecast cycle. Policymakers review these downside scenarios before approving new infrastructure.

TEST 3 — AFTER THE DEAL IS DONE.

Does the Facility Leave the Public Better Off After Its Contracts Expire?

The third test focuses on the 20- to 30-year impacts of infrastructure and energy decisions that outlast contracts, developers, and the regulatory approvals that enabled them. Passing the test requires all three policy elements: new electricity demand served by genuinely new clean supply rather than fossil-fuel lock-in, a transparent siting decision made with the affected community heard, and public subsidies tied to the delivery of real benefits.

ELEMENT 7:
New-and-Matched Clean Energy

What it targets

This element addresses the risk of hyperscale load locking in climate pollution and other health and environmental harms, a risk not captured by the scenario matrix's focus on the direct financial risks of who pays and whether electricity demand actually materializes. That matrix tracks cost responsibility and demand materialization, but not the type of infrastructure built or its long-term impact. Even if demand fully materializes and hyperscalers bear the costs, relying on fossil fuels can still lock in climate pollution, stifle clean energy investment, and concentrate pollution in vulnerable communities. Element 7 prevents the "Hyperscaler Bears Costs" column from becoming a path to more fossil fuel use by requiring that new infrastructure for large loads is clean. It works with Element 4: time-matching requirements and real-cost pricing together create a direct incentive to source clean energy when and where it's most affordable.

What it looks like in practice

A new 250 MW campus must source a defined minimum percentage of its annual electricity from new, additive clean energy as an enforceable condition of the interconnection agreement, with financial penalties that scale with the size and duration of any shortfall. Renewable Energy Credits (RECs) and power purchase agreements with existing clean energy facilities *do not satisfy this policy element* because they pull clean energy attributes from the grid to attribute to the hyperscale project, with no requirement that those attributes be replaced by new generation resources.

Eligible pathways are limited to behind-the-meter clean generation, power purchase agreements for projects entering commercial operation after a defined date, or directly interconnected new renewable facilities. Where a developer claims additive supply is physically or technically infeasible, a narrow variance process is available; cost alone is not grounds for exemption, and any variance is time-limited and publicly disclosed. Once defined metering and market-infrastructure thresholds are met, the standard shifts from annual average to hourly matching.

Core legislative requirements

- **A minimum clean energy threshold.** Large facilities must source a defined minimum percentage of their energy from clean energy as an enforceable condition of interconnection. Without this floor, hyperscale demand growth becomes a default driver of new fossil fuel generation.
- **Renewable energy credits alone do not count.** Meeting the clean energy requirement by purchasing paper credits that do not represent new generation added to the grid is explicitly prohibited. Compliance requires genuinely additional clean energy supply – new generation that would not have been built without the facility's demand, and that can be traced to a specific source delivering power to the same grid.
- **Phased milestones with defined consequences for missing them.** The requirement is structured as a phased timeline with hard deadlines. For example, 50% clean energy by year three, 80% by year seven, and 100% by year ten. A facility that misses a milestone is subject to defined financial penalties that increase with the size and duration of the shortfall.
- **A narrow, time-limited exemption for genuinely constrained sites.** Where a developer can demonstrate, through independent verification, that sourcing additional clean energy at a specific location is physically or technically infeasible, a defined variance process is available. Financial cost alone does not qualify. Any exemption is time-limited, publicly disclosed, and subject to mandatory review if conditions change.
- **Hourly matching with a defined trigger for when it takes effect.** Once metering and market-infrastructure thresholds are met, the requirement shifts from annual to hourly matching, so clean energy is sourced during the same hours it is consumed rather than on an annual average.

ELEMENT 8:
Community-Protective Siting Standards

What it targets

This element addresses a risk not captured by the scenario matrix's main focus, but one that affects the credibility of every outcome – whether cost and responsibility decisions are made openly and with meaningful public input. By requiring the same procedural rigor as major grid projects (e.g., independent studies, risk assessments, site alternatives, and a set public comment period), this element ensures decisions are both open and inclusive, and guarantees that communities most affected by hyperscale development have a formal opportunity to be heard before any decision is finalized.

What it looks like in practice

A developer proposing a 200 MW or greater campus must complete three studies before the PUC will consider an interconnection application:

1. A grid impact study (what the connection requires and costs);
2. A ratepayer risk assessment (what ordinary customers face if the facility underperforms or shuts down);
3. And a site alternatives analysis (proof that other locations were seriously considered).

The alternatives analysis must directly address community impact: if the proposed site is near a community already facing high pollution, the developer must show that less-harmful alternatives were considered and explain why they were not chosen. The process concludes with an accessible public comment period, after which the PUC issues a formal approval decision on the record.

Core legislative requirements

- **Grid impact studies must be completed before any large connection is approved.** For any facility above a defined threshold, an independent study establishes what the connection will require — transmission, substations, generation capacity — and what it will cost.
- **A consumer risk assessment alongside the grid impact study.** A separate assessment models the financial exposure ordinary customers face under downside scenarios (partial materialization, facility exit, stranded assets) before any commitment is made. Policymakers cannot make an informed decision without knowing what they are asking consumers to absorb.
- **A site alternatives analysis before the location is locked in.** The developer must demonstrate that alternatives were genuinely considered, such as locations with available grid capacity, where clean energy supply is more accessible, or where community impacts are lower. Approving a location without this analysis eliminates options that may have been significantly better for the grid, for consumers, and for surrounding communities.
- **A public comment period and a defined PUC approval step.** Before an interconnection agreement above the defined threshold takes effect, it must pass through a formal public comment period and receive explicit PUC approval, so affected communities, ratepayers, and advocates have a defined opportunity to be heard before the decision is final.

ELEMENT 9:
Performance-Tied Incentives

What it targets

Addresses the risk of unconditional subsidies by making public incentives dependent on facilities delivering their promised demand. Incentives are not given upfront and left in place regardless of performance – they are tied to actual delivery. In the scenario matrix, this focuses on the partial- and non-materialization rows: if a facility underperforms or exits, it loses the fiscal benefits meant for full operation. Linking incentive value to real demand moves outcomes away from “Consumers Bear Costs.” Public support is clawed back before it adds to the risk of stranded assets.

What it looks like in practice

A campus receives a 10-year property tax abatement conditional on maintaining at least 80% of its contracted load commitment throughout the abatement period. If the facility exits or drops below that threshold in year six, a clawback provision requires repayment of a pro-rated share of the abatement value.

Core legislative requirements

- **Tax incentives expire if load commitments are not met.** Any abatement or incentive granted to a large facility is time-bound and tied to defined load commitments; failing to meet them causes the incentive to step down or expire automatically.
- **Developers must repay incentives if they exit early.** If a facility that received public incentives exits, significantly reduces its load, or fails to meet defined operational milestones within a specified period, clawback provisions require repayment of some or all of the support received.

Existing Facilities — A Different Legislative Problem

The nine elements are framed as conditions on new interconnection agreements. That is the right starting point: the point of entry is where requirements are easiest to establish and hardest to circumvent. But a significant and growing share of ratepayer exposure comes from facilities already operating under legacy agreements that predate these requirements entirely. New rules that apply only to new facilities leave the largest, most concentrated loads on the system under a weaker set of obligations. Closing that gap requires a different legislative approach.

The core problem: Legacy agreements create a double standard: new facilities must meet requirements that older ones do not, even though existing facilities are often the largest and most concentrated users on the grid

Three mechanisms are available (see Table 3), listed from least to most disruptive. Begin with the option that best fits your political environment.

Applying Requirements to Legacy Agreements

Mechanism	When to use it	What it requires	The defensible argument
Update the rules at contract renewal	A legacy agreement is expiring or coming up for renegotiation	The facility must meet current standards as a condition of renewal (i.e., financial commitments, demand response enrolment, and clean energy requirements) from the new agreement date forward	The agreement was always going to be renegotiated. This simply ensures the new terms reflect what is now required of everyone else
Apply current rules when a facility asks for something new	An existing facility wants to add load, requests new infrastructure, or applies to renew an incentive	Current standards apply to the change being requested; for example, a facility asking for a 50 MW expansion is in the same position as a new applicant for that increment of capacity	The facility is already in front of the regulator seeking approval
Set a compliance deadline for all large existing facilities	Applied across all facilities above a defined size threshold, regardless of whether they are seeking anything new	A fixed schedule by which all large existing facilities must meet demand response, financial assurance, and clean energy requirements, with defined financial penalties for missing milestones	A reliability event or demonstrated ratepayer harm has created the political conditions for direct intervention. This is analogous to air quality compliance schedules for industrial emitters

Enforcement — A Cross-Cutting Requirement

Every element in this framework depends on enforcement to function. A take-or-pay contract with no penalty for non-payment is not meaningfully different from a voluntary commitment. A demand response enrolment requirement with no consequences for failing to curtail does not protect the grid. Before any element can be considered fully implemented, four enforcement questions must be answered in the enabling legislation or PUC rulemaking.

- **What is the verification mechanism?** Who confirms that a developer is meeting its commitments, on what schedule, and using what data sources? Independent verification — by the PUC, a designated third party, or an independent office that is provided with adequate budget and staffing to carry out this responsibility— is significantly more reliable than self-reporting.
- **What are the penalties for non-compliance?** Penalties must be calibrated to exceed the cost of non-compliance — otherwise rational actors will simply pay them. A penalty for demand-response failure, for example, should exceed the value of the load the facility failed to curtail.
- **What remedies are available to ratepayers when requirements are violated?** Statutory provisions that give ratepayer advocates standing to intervene in enforcement proceedings are more durable than agency discretion alone.
- **What is the appeals process, and does it allow developers to delay compliance while appeals are pending?** Open-ended appeals processes are a common mechanism for effectively nullifying requirements that are valid on their face.

These questions should be answered explicitly in bill drafting. Do not assume that existing enforcement mechanisms in utility regulation are adequate — most were designed for investor-owned utilities operating under long-term rate cases, not for the rapid-cycle, developer-driven dynamic that characterizes hyperscale interconnection.

CASE STUDY: EVALUATING A LEGISLATIVE PORTFOLIO— VIRGINIA'S 2025 - 2026 DATA CENTER DOCKET

Virginia offers a useful test of the framework because it combines the world's largest concentration of data center development with one of the nation's most active legislative responses. The facts are drawn from the public record, including the Joint Legislative Audit and Review Commission (JLARC) report on data centers, legislative records from the 2025 and 2026 General Assembly sessions, and news reports from that time. The framework's assessment, however, is interpretive. Rather than evaluating individual bills in isolation, it asks whether Virginia's legislative portfolio, taken as a whole, addressed the principal governance challenges associated with hyperscale data center growth.

The answer is mixed. Across two sessions, Virginia produced an unusually large volume of legislation and, in 2026, enacted several measures that failed a year earlier. Yet significant policy gaps remained across every stage of the project life cycle. Applied across the 2025 and 2026 sessions, the framework's three tests yielded three findings:

- **Before approval, the stage that matters most, coverage advanced, but the portfolio fell short.** In 2025, the bills that would have verified demand, placed developer capital at risk, or assigned infrastructure costs to the customers driving them nearly all failed. In 2026, legislation establishing the principle of cost allocation for high-load customers was enacted, although subsequent amendments replaced mandatory cost-shifting provisions with more discretionary language, and independent demand verification remained optional.
- **During operation, safeguards remained only partially developed across both sessions.** The measures that passed took the form of studies, pilots, and procedural reforms rather than operational requirements. In both 2025 and 2026, Virginia did not enact real-cost pricing, impose mandatory curtailment before households, or conduct downside stress testing tied to large new loads, even as regional reliability pressures intensified.
- **Over the long term, protections expanded but stopped short.** Legislation establishing community siting requirements and water-use transparency, which did not advance in 2025, was enacted in 2026, expanding the state's oversight of data center development. Yet the state's multi-billion-dollar tax incentive remained largely disconnected from performance, and new demand was still not matched to new clean generation.

WHY VIRGINIA

Virginia's experience illustrates why data center governance has become a statewide policy issue rather than simply a land-use question. Northern Virginia hosts approximately 4,140 MWs of operational data center capacity, more than twice that of the next-largest global market and roughly 13% of worldwide capacity. As development accelerated, policymakers increasingly confronted questions that extend well beyond zoning: Who should pay for new transmission infrastructure? How should utilities plan for unprecedented demand growth? What obligations should accompany billions of dollars in public tax incentives?

JLARC's December 2024 report sharpened those questions. It concluded that unconstrained data center expansion could double Virginia's electricity demand within a decade, requiring approximately 150% more electric generation and 40% more transmission capacity. It also emphasized an important distinction. Existing electric rates generally assign costs appropriately, but future system expansion would raise costs for all customers unless regulatory policy evolved alongside demand growth. At the same time, the report highlighted substantial uncertainty surrounding long-term forecasts, noting that improvements in processor efficiency and cooling technologies could materially reduce future electricity demand.

Those findings created an unusual legislative environment. More than three dozen data center and related energy bills were introduced during the 2025 session, spanning utility regulation, cost allocation, electric system planning, siting, environmental review, transparency, and tax policy. The volume of activity reflected broad bipartisan recognition that existing policy frameworks were no longer sufficient. The more difficult question was whether those proposals collectively formed a coherent governance strategy.

Reading the Docket as a Portfolio

Evaluating legislation bill by bill suggests a busy and productive session. Evaluating the same legislation as a portfolio produces a different picture.

The framework organizes proposals around three questions corresponding to the life cycle of a data center project: What occurs before approval? What governs operations after interconnection? What obligations remain over the project's long-term life? Applying that structure to Virginia's 2025 docket reveals a striking imbalance.

Before Approval: The Ambitious Bills That Died

The strongest legislative attention focused on the period before project approval. Lawmakers introduced bills addressing cost allocation, developer responsibility, statewide review, and operational transparency. Several proposals would have directed the State Corporation Commission to determine whether non-data-center customers were subsidizing infrastructure built primarily for hyperscale facilities. Others sought new reporting requirements or centralized review of major energy facilities.

Few of these measures ultimately became law. The principal cost-allocation legislation passed the Senate but failed in conference committee. Statewide review proposals died in committee. Transparency requirements likewise failed to advance. The only surviving measure concerning customer classification was substantially narrowed during the legislative process, leaving regulators to study existing authority rather than establishing new requirements.

As a result, the enacted portfolio did not require independently verified demand forecasts, meaningful developer financial commitments before infrastructure investment, or comprehensive cause-and-pay cost allocation. From the framework's perspective, the approval stage remained largely unaddressed despite substantial legislative effort.

While It Runs: The Gap After the Grid Connection

The most significant gap emerged after facilities connect to the grid.

Several enacted bills examined demand response and utility planning, but nearly all took the form of studies, pilot programs, or procedural reforms rather than operational requirements. Legislation directed the Department of Energy to evaluate demand response programs and authorized a voluntary virtual power plant pilot. Integrated resource planning requirements were revised, but without requiring utilities to evaluate downside demand scenarios or reassess investments when very large new loads materialize.

Similarly, proposals concerning electricity pricing never advanced beyond preliminary consideration. Virginia enacted neither locational pricing nor time-varying rates designed specifically for hyperscale facilities.

This distinction between studies and requirements is central to the framework. Studies improve information. Pilot programs test implementation. Neither, however, changes operating incentives. Consequently, the framework characterizes Virginia's operational safeguards as only partially developed.

After the Deal: Long-Term Governance Remained Incomplete

The legislative record becomes more uneven when viewed over the decades-long life of data center infrastructure.

The session's most prominent siting proposal would have allowed local governments to require assessments examining impacts on water resources, agriculture, forests, parks, and nearby communities before approving new facilities. Although the legislation passed both chambers, it was returned by the governor with amendments that substantially weakened its provisions and ultimately did not become law.

Likewise, proposals linking Virginia's data center sales-and-use tax exemption to clean-energy or efficiency requirements failed in committee. More general clean energy legislation advanced during the session, including measures supporting renewable generation, storage, and solar deployment. None, however, required new data center demand to be matched with additional clean electricity resources or conditioned public incentives on measurable performance.

As a result, the framework finds that Virginia made limited progress toward integrating long-term environmental, fiscal, and community protections into its governance model.

WHAT WAS NEVER FILED

Much of the framework's value lies in identifying which issues received comparatively little legislative attention. Portfolio analysis offers something individual bill analysis cannot: it reveals not just what legislators debated, but which governance functions remained largely outside that conversation.

No significant proposal addressed existing hyperscale facilities operating under legacy agreements, even though those facilities represent the largest concentration of electric demand in the Commonwealth. Nearly every bill applied prospectively, to new projects only.

Enforcement received similarly little attention. Most enacted legislation established studies, reporting requirements, or future regulatory proceedings, but few provisions specified independent verification, meaningful penalties, opportunities for affected ratepayers to participate, or mechanisms to limit procedural delay. These implementation details often determine whether statutory requirements change behavior in practice.

Revisiting the Portfolio: The 2026 Session

The 2026 legislative session enabled an evaluation of whether the challenges in Virginia's 2025 data center agenda stemmed from temporary political factors or deeper structural issues. Rather than retreating from the issue, lawmakers expanded their focus significantly. The General Assembly reviewed 61 data center-related bills, passed 15, and deferred 46 to the 2027 session. Several proposals that had previously failed in 2025, such as those on electricity cost sharing, community siting, and water-use transparency were finally passed. However, when assessed within the project life cycle framework, the policy portfolio remained inconsistent. Virginia enhanced its regulatory approach but still left notable governance gaps unaddressed.

The most notable development before project approval was the enactment of legislation addressing electricity cost allocation. House Bill 1393 and Senate Bill 253 revived the central objective of 2025 proposals by seeking to ensure that the costs associated with serving large data center loads would not be borne by residential customers. During the legislative process, however, the bills were substantially revised. The explicit requirement to assign infrastructure and capacity costs to high-load customers was replaced with language directing the State Corporation Commission to consider avoiding cost shifts to other customers, while the scope of the legislation was narrowed to existing

utility customers. The legislation enacted an important policy principle, but in a narrower and discretionary form than originally proposed.

Operational oversight also expanded, although largely through voluntary rather than mandatory mechanisms. The General Assembly directed electric utilities to develop demand-flexibility programs for customers with loads of 25 MWs or greater, creating a pathway for large users to reduce consumption during periods of system stress. Participation, however, remained voluntary, and lawmakers did not adopt mandatory curtailment requirements or priority protections for residential consumers. The legislature also failed to require utilities to incorporate downside demand scenarios or independent verification into long-term planning. These omissions grew more consequential as regional reliability concerns intensified. During the session, PJM failed to meet its reserve-margin target, and its independent market monitor recommended pausing the interconnection of additional large data center loads until reliability concerns were addressed. The developments underscored the growing importance of planning tools that remained absent from Virginia's governance framework.

The post-approval phase saw the most substantial expansion of community protections. House Bill 153 established a permitting process for data centers with loads of 100 MWs or more, requiring applicants to evaluate potential impacts on surrounding communities, including noise, and House Bill 496 introduced monthly reporting requirements for water supplied to data centers. Additional legislation required cleaner Tier IV backup generators, strengthening the local environmental protections tied to facility operations.

Fiscal policy, however, remained the dominant issue of the session. Debate centered on Virginia's sales tax exemption for qualifying data centers, estimated to cost between approximately $1.6 billion and $1.9 billion annually. Competing proposals to repeal, preserve, or modify the incentive became a central point of budget negotiations, contributing to a budget impasse that extended for more than three months. The eventual compromise preserved the exemption while creating a temporary energy tax projected to generate approximately $600 million annually. Although the agreement increased public revenues from the sector, it left the underlying incentive structure largely unchanged. The exemption remained largely disconnected from measurable performance outcomes such as clean energy procurement, infrastructure investment, or community benefits, and no clawback provisions were adopted.

Just as significant were the reforms that failed to advance. The legislature again declined to require independent verification of electricity demand forecasts, mandatory demand curtailment during system emergencies, scenario-based planning for rapid load growth, or requirements that new electricity demand be matched with clean generation. Existing data centers continued to receive relatively limited regulatory scrutiny, while public incentives remained largely unconnected to ongoing performance requirements. Many of these proposals were carried over to the 2027 session, suggesting that Virginia's policy agenda continued to expand while the most politically difficult elements of comprehensive governance remained unresolved.

Viewed across the project life cycle, the 2026 session reflects a pattern of incremental institutional progress rather than comprehensive reform. The state strengthened transparency, community oversight, and recognition of cost allocation concerns, but many enacted measures were narrowed during the legislative process, and several foundational safeguards continued to rely on voluntary participation or future legislative action. Virginia's experience therefore illustrates how policy portfolios can

become broader over time without yet achieving comprehensive coverage of the risks associated with large-scale data center development.

Lessons for Policymakers

Virginia's experience demonstrates that legislative volume should not be mistaken for comprehensive policy coverage. More than three dozen bills in 2025 and more than sixty in 2026 addressed data center development. Yet across both sessions, the portfolio advanced unevenly, leaving important governance functions only partially addressed or absent, even as the volume of activity and the urgency behind it kept rising.

The broader implication extends beyond Virginia. As additional states confront rapid growth in hyperscale data centers, legislatures are likely to experience similar surges of overlapping proposals introduced across multiple committees and policy domains. Evaluated individually, many of those bills will appear reasonable. Evaluated collectively, they may still leave significant gaps.

The framework presented in this paper is intended to make those gaps visible. Organizing legislation by the life cycle of a project, rather than by committee jurisdiction or statutory topic, provides policymakers with a systematic method for evaluating whether a legislative portfolio functions as a coherent governance strategy rather than simply as a collection of individual bills.

Glossary

Adaptive planning: A regulatory requirement that utilities update their Integrated Resource Plans more frequently than the standard three-to-five-year cycle, or upon the occurrence of defined trigger events such as a large-load addition above a threshold MW level. The adaptive planning requirement in Element 6 is designed to prevent utilities from committing to long-lived assets based on demand forecasts that may be outdated within 18 months of filing.

Additionality: The requirement that clean energy credited toward a compliance obligation must come from generation capacity that would not otherwise have been built. A project is additive if it has a commercial operation date after the compliance obligation was established, is physically deliverable to the load zone, and has not been counted toward another customer's compliance obligation. Additionality is the defining design feature of Element 7. Without it, a clean energy requirement can be satisfied by purchasing accounting credits from existing generation rather than driving new supply.

Behind-the-meter generation: Generation capacity located on the customer's side of the utility meter, not connected to the bulk power grid. Behind-the-meter generation reduces the facility's net grid consumption but is not directly visible to the grid operator. Solar panels on a data center roof are behind-the-meter generation. Behind-the-meter resources can qualify toward additionality requirements in Element 7 if they meet the commercial operation date threshold and additionality test.

Bond / Letter of credit: A financial instrument posted by a developer before dedicated infrastructure is built, equal to a defined share of that infrastructure cost. The bond is the primary enforcement mechanism for Element 2. If the developer exits or fails to meet take-or-pay obligations, stranded asset recovery flows from the bond before costs reach the general rate base. A letter of credit functions similarly but is issued by a financial institution rather than posted directly by the developer. The methodology for determining bond adequacy must be set by statute or PUC rulemaking.

Carbon lock-in: The long-term commitment to carbon-emitting infrastructure that results from building fossil fuel generation capacity to serve new demand. Because gas-fired plants operate for decades, approval decisions made today constrain decarbonization options long after the contracts that justified them expire — the 20- to 30-year horizon that Tier 3 of the framework addresses. Carbon lock-in is the climate dimension of the gas dependency risk, which Element 7 addresses but which cost-only frameworks leave unmanaged.

Clawback provision: A contractual or statutory provision requiring a subsidy recipient to return some or all of the benefit received if specific conditions, such as load commitments, employment thresholds, or clean energy requirements, are not met. In Element 9, clawbacks are calculated from the date of non-compliance rather than the original agreement date, so a facility that operates at full commitment for several years and then exits repays a pro-rated share of the abatement value rather than the full package. Clawback provisions are the primary enforcement mechanism for incentive conditionality.

Commercial operation date: The date on which a generating facility begins commercial operation. Used in additionality requirements to establish whether a project qualifies as new generation or existing generation that does not satisfy the additionality requirement. Bill drafters must define the

commercial operation date threshold precisely; an imprecise threshold creates a gap through which developers can satisfy additionality requirements using generation that predates the policy intent.

Cost responsibility agreement: A contract assigning the cost of dedicated infrastructure — substations, transmission lines, or generation capacity built primarily to serve a single customer — to that customer before any portion is eligible for rate base recovery. Cost responsibility agreements are a core legislative requirement of Elements 2 and 3. Their enforceability depends on a clear statutory definition of what constitutes "dedicated" infrastructure — the most frequently litigated question in large-load interconnection cost allocation.

Curtailment: The reduction of electricity consumption in response to a grid operator signal. Mandatory curtailment under Element 5 requires enrolled facilities to reduce load by a specified amount within a defined response time during grid emergencies, ahead of any residential curtailment. Curtailment protocol design, specifying which loads within a facility can be curtailed without affecting critical operations, is technically complex and requires utility and grid operator cooperation before the legislative requirement is enacted.

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 use, creating stranded asset risk. The legal and regulatory definition of what constitutes "primarily serving" a single customer is the most consequential and most frequently litigated question in hyperscale interconnection cost allocation. Elements 2, 3, and 8 all depend on a workable statutory definition of this term.

Demand forecast: A projection of the electricity load a proposed facility is expected to place on the grid over a defined period. Developers who supply these forecasts have a direct financial stake in the approval of the infrastructure they justify. Element 1 therefore requires forecasts to be audited against comparable facility data, independent engineering analysis, and conservative downside scenarios before they are incorporated into an IRP or interconnection agreement. The reliability of every subsequent element depends on the quality of that verified baseline.

Demand response: A program under which electricity customers agree to reduce their consumption in response to grid operator signals, typically during peak demand periods or grid emergencies. Mandatory demand response enrollment under Element 5 converts hyperscale facilities from passive contributors to grid stress into active reliability assets. In ERCOT, demand response enrolment should be treated as a current-session priority given tight reserve margins.

ERCOT (Electric Reliability Council of Texas): The grid operator for most of Texas, running an energy-only wholesale market with nodal pricing and recurring reserve-margin pressure. Because large ERCOT loads often procure power through retail electric providers on fixed contracts that obscure nodal price signals, Element 4 recommends requiring REPs serving loads above a defined threshold to offer nodal-price-indexed contracts, and Element 5 treats demand response enrolment as a current-session priority.

Financial assurance: The bond, letter of credit, or equivalent instrument required under Element 2 before a utility breaks ground on dedicated infrastructure. Financial assurance is the primary Tier 1 mechanism for ensuring that stranded asset costs fall on the developer before reaching the general rate base. Its enforceability depends on: (1) a PUC-approved methodology for setting bond amounts;

(2) third-party confirmation of bond value; and (3) statutory standing for ratepayer advocates to challenge bond adequacy determinations.

Forecast opacity: The condition in which demand projections for a proposed facility rely primarily on the developer's own internal load estimates rather than independent analysis, making it difficult for regulators to assess the credibility of the forecast or its downside scenarios. Forecast opacity is one of the eight risks identified in Taming Data Center Turmoil: Hyperscalers Push Electricity Markets to Their Limits and is addressed primarily by Element 1, with Element 6 in support (Table 1).

Gas dependency: The risk that meeting hyperscale demand requires significant new natural gas capacity, creating a decades-long commitment to fuel-cost volatility, carbon emissions, and air-quality impacts on host communities. Identified in Table 1 as "gas dependency and fossil fuel facilitation," this risk is addressed primarily by Element 7, with Elements 4 and 8 providing supporting mechanisms.

Grid impact study: An independent study, completed before an interconnection agreement is approved, establishing what a proposed connection will actually require from the grid (e.g., new transmission, upgraded substations, additional generation capacity) and what it will cost to deliver. Impact studies are required under Element 1 for facilities above a defined size threshold and under Element 8 as one of three studies that must precede PUC consideration of an interconnection application.

Hyperscale data center (hyperscaler): A data center campus, typically built by or for major cloud and artificial intelligence providers, whose electricity demand is measured in the hundreds of megawatts. Hyperscale facilities are distinguished less by size alone than by how they arrive: in months rather than years, clustered in specific locations rather than spread across the system, under contracts that expire before the infrastructure built to serve them does, and backed by developers with a direct financial stake in the demand forecasts that justify that infrastructure.

Independent verification: Confirmation by an entity outside the developer's organization of the claims the framework depends on. Under Element 1, a designated independent body audits large-load demand forecasts against engineering feasibility, benchmarks from comparable facilities, and conservative downside scenarios before any infrastructure commitment is approved, adding an estimated 60–120 days to the approval process. Under the framework's enforcement requirements, independent verification of ongoing compliance is significantly more reliable than self-reporting.

Interconnection agreement: The contract between a utility or grid operator and a customer establishing the terms under which the customer connects to the grid. The interconnection agreement is the framework's primary vehicle for attaching requirements — take-or-pay clauses, minimum use commitments, clean energy milestones, curtailment obligations — because the point of entry is where requirements are easiest to establish and hardest to circumvent.

IRP (Integrated Resource Plan): A utility's long-term plan for meeting electricity demand through a combination of generation, transmission, distribution, and demand-side resources. IRPs are typically filed with state PUCs every three to five years and used to justify infrastructure investment decisions. Element 6 requires accelerated IRP update cycles — triggered by large-load additions exceeding a defined threshold — and mandatory downside-scenario analysis as a condition of PUC approval for new generation or transmission commitments.

Large-load rate class: A separate rate classification for customers above a defined MW threshold, segregating hyperscale customers from general service tariffs. A large-load rate class is a core

legislative requirement of Element 2. Without it, take-or-pay and financial assurance provisions cannot be applied to hyperscale customers without applying them to all large customers on the same tariff. The MW threshold for inclusion should be set at a level that captures hyperscale facilities while excluding traditional large industrial customers who are not driving the infrastructure risk.

Legacy agreement: An interconnection agreement or incentive package executed before current requirements took effect. Because the nine elements attach at the point of new interconnection, facilities operating under legacy agreements — often the largest and most concentrated loads on the system — would otherwise face a weaker set of obligations than new entrants. Table 3 presents three mechanisms for closing that gap: updating terms at contract renewal, applying current rules when a facility requests something new, and setting a compliance deadline for all large existing facilities.

Load commitment: A binding obligation by a large electricity customer to consume a specified minimum level of power over a defined period. Load commitments are the mechanism through which clawback provisions in Element 9 are triggered — a facility that drops below its committed load level activates the clawback calculation — and the basis for the minimum use commitments required in interconnection agreements under Element 3. Load commitment definitions must specify the minimum threshold, the measurement period, and the methodology for calculating compliance. Vague load commitments are unenforceable load commitments.

Nodal pricing: A pricing methodology that assigns different electricity prices to different nodes — points of injection or withdrawal — on the grid, based on local supply and demand conditions and transmission constraints. Nodal pricing creates the locational signals that incentivize hyperscale facilities to site where grid capacity exists rather than where real estate is cheapest. Wholesale nodal pricing already exists in most RTO states; the policy gap Element 4 addresses is retail translation, ensuring that large loads actually face those signals. In vertically integrated states, nodal pricing requires fundamental rate design reform, making this the most technically complex element of the framework.

Pre-application study: An engineering assessment funded by the developer and completed before any interconnection agreement is signed, assessing grid impacts, required upgrades, and estimated costs. Pre-application studies are a core requirement of Element 3, typically adding 30–90 days to the approval process. A developer unwilling to fund a study assessing the costs their project will impose is a developer whose project the utility and ratepayers should be cautious about sponsoring.

PUC (Public Utilities Commission): The state regulatory agency responsible for overseeing investor-owned utilities, approving rate cases, and setting the terms of interconnection agreements. The specific name varies by state — Public Service Commission, Corporation Commission, and other variants are used in different jurisdictions. PUCs operate under authority granted by state legislatures; the framework elements in this document require statutory direction to the PUC rather than PUC action alone.

Rate base: The total value of utility assets on which a utility is permitted to earn a regulated rate of return. Infrastructure investments that enter the rate base are recovered from all ratepayers through electricity rates. The core ratepayer protection goal of Elements 2 and 3 is to ensure that dedicated infrastructure costs flow to the driving customer before entering the rate base — preventing the socialization of costs that ratepayers neither caused nor benefit from.

Ratepayer risk assessment: A study, required under Element 8 alongside the grid impact study, that models the financial exposure ordinary customers face under downside scenarios — partial

materialization, facility exit, stranded assets — before any commitment is made. The grid impact study establishes what a connection will cost; the ratepayer risk assessment establishes who bears that cost if things go wrong. Policymakers cannot make an informed approval decision without knowing what they are asking ratepayers to absorb.

REC (Renewable Energy Certificate): A tradeable instrument representing the environmental attributes of one megawatt-hour of renewable electricity generation. RECs can be bought and sold separately from the electricity itself. Unbundled RECs — purchased separately from the power — do not guarantee that the buyer is consuming renewable electricity; they transfer accounting credits only. Element 7 explicitly prohibits unbundled RECs from satisfying clean energy compliance requirements. Bill drafters must be precise: bundled PPAs with new projects may qualify if the project meets the commercial operation date threshold.

Retail electric provider (REP): In retail-choice markets such as ERCOT, the competitive supplier from which customers purchase electricity. Fixed-price REP contracts can insulate large loads from the locational and time-based price signals that exist at the wholesale level. Element 4 addresses this gap by recommending that REPs serving loads above a defined threshold be required to offer nodal-price-indexed contracts.

RTO/ISO (Regional Transmission Organization / Independent System Operator): An organization that coordinates, controls, and monitors a multi-state electricity grid. RTOs operate wholesale electricity markets and manage transmission access under federal jurisdiction, while states retain authority over distribution, retail rates, and siting. Elements 3, 4, 5, and 6 all carry RTO-state notes in their tradeoff tables because the boundary between state and federal jurisdiction is the most important constraint on bill drafting in competitive wholesale states.

Scenario matrix: The three-by-three matrix, reproduced as Table 2 from Taming Data Center Turmoil: The Rise of Hyperscale Load, that crosses demand outcomes (full materialization, partial materialization, non-materialization) with cost responsibility (hyperscaler bears costs, shared costs, ratepayers bear costs). Each framework element identifies the cells it targets; the framework's overall purpose is to keep outcomes out of the "Ratepayers Bear Costs" column regardless of which demand row materializes.

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 the key indicator of single-customer concentration and a primary consideration in the forecast review required under Element 1.

Single-customer concentration: The condition in which a single customer accounts for a disproportionate share of a utility's peak demand within its service territory. High concentration creates vulnerability: if that customer exits or reduces load, fixed cost recovery is destabilized for all remaining ratepayers. Single-customer concentration is one of the eight risks the framework addresses; Element 2 is the primary response, supported by Elements 1 and 8 (Table 1).

Site alternatives analysis: A study, required under Element 8 before a specific site is approved, demonstrating that alternative locations were genuinely considered — including sites where the grid has available capacity, where clean energy supply is more accessible, or where community impacts are lower. Where the proposed site is near a community already facing high pollution, the analysis must show that less harmful alternatives were evaluated and explain why they were not chosen.

Stranded asset: Infrastructure that has been financed and built but earns less revenue than projected — typically because the customer it was built to serve has exited, reduced consumption, or failed to materialize as expected. Stranded asset costs are recovered from ratepayers through rate increases unless cost responsibility agreements and financial assurance provisions require the departing customer to bear them. The stranded asset (partial) and stranded asset (full) cells of the scenario matrix (Table 2) represent the outcomes that Elements 1, 2, and 3 are most directly designed to prevent.

Take-or-pay contract: A contract requiring the buyer to pay for a contracted quantity of capacity or energy, whether or not it is actually consumed. In the context of large-load interconnection, take-or-pay provisions under Element 2 require developers to pay for their contracted grid capacity even if actual load falls short of projections. Take-or-pay contracts are most enforceable when the contracted capacity figure is grounded in independently verified demand baselines, which is why Element 1 is the precondition for Element 2 to function at its full protective value.

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, an outcome that has damaged the credibility of data center incentive programs in several states. Element 9 requires abatements to be time-bounded and conditioned on load commitments, with clawback provisions if commitments are not met.

Time matching: A clean energy compliance standard requiring clean energy to be delivered at the same time it is consumed, rather than averaged over an annual period. Hourly time matching — the most demanding standard — requires clean supply for every hour of operation. Annual average matching allows a facility to consume gas or coal power at peak and offset it with overnight solar credits; hourly matching closes that gap. Element 7 phases hourly matching in once defined metering and market infrastructure thresholds are met, rather than requiring it from the first compliance year.

Trigger threshold: A defined size level, typically in megawatts, at which specific regulatory requirements automatically apply. Trigger thresholds appear throughout the framework: the size threshold above which impact studies are required under Elements 1 and 8, the rate-class and financial-assurance thresholds under Element 2, the mandatory demand response enrolment threshold under Element 5, and the large-load-addition trigger for accelerated IRP filings under Element 6. Setting thresholds too high allows proposals to bypass requirements; setting them too low creates administrative burden without proportionate protective benefit.

Variance process: A defined regulatory pathway through which a developer can demonstrate that compliance with a specific element requirement is infeasible at their proposed site, and obtain a modification or exemption from that requirement. Element 7 requires a variance process for sites where additive clean energy supply is demonstrably infeasible; financial cost alone does not qualify, and any variance granted is time-limited, publicly disclosed, and subject to mandatory review if conditions change. The variance process must require the developer to demonstrate infeasibility — the default must be compliance, not exemption. A variance process that routinely exempts constrained sites without requiring the developer to show their work is not a variance process; it is a default exemption dressed as flexibility.