From Trilemma to Triumph

Powering California's Data Center Boom

By Liang Min, Rajanie Prabha, Ram Rajagopal, and William Chueh

ELECTRIC POWER SYSTEMS WORLDWIDE FACE a persistent trilemma: balancing reliability, affordability, and sustainability. In California, this challenge is intensifying as electricity demand surges alongside the rapid growth in artificial intelligence (AI) and data centers. The state already has some of the highest electricity rates in the United States, one of the largest renewable and storage portfolios, and growing exposure to wildfire and climate risks. Meeting this wave of demand without driving up costs, compromising reliability, or increasing emissions will test the limits of California's grid and the creativity of its operators and regulators.

This article outlines how California can turn AI-driven demand from a perceived burden into a catalyst for a cleaner, more resilient, and more affordable grid. Drawing on system data, we identify three complementary strategies: 1) unlock underutilized transmission capacity to power AI-driven load growth, 2) target clean and dispatchable resources to areas of highest grid impact, and 3) implement equitable cost-recovery mechanisms so data centers fund needed upgrades without burdening other customers.

If designed and implemented effectively, these strategies turn data centers into partners in modernizing California's electric grid and

Liang Min (liangmin@stanford.edu), Rajanie Prabha (rajanie@stanford.edu), Ram Rajagopal (ramr@stanford.edu), and William Chueh (wchueh@stanford.edu) are with Stanford University, Stanford, CA 94305-2004 USA. Liang Min is the corresponding author.

Digital Object Identifier 10.1109/ESM.2026.3681527
Date of current version: 29 July 2026

spread infrastructure costs across more electricity users. The result is a practical pathway to turn today's trilemma into tomorrow's triumph for California and a model that other regions can adapt as digital demand scales.

California's Data Center Boom

The AI era is reshaping electricity demand at a pace the United States has not seen since the postwar industrial boom. The California Energy Commission (CEC) projects California Independent System Operator (CAISO) peak demand rising from 48.3 GW in 2024 to roughly 68 GW by 2040. This increase is equivalent to powering about 20 million additional homes.

Local utilities are already feeling the effects. Silicon Valley Power's system peak is about 714 MW with a 76.6% load factor, and roughly 91.7% of its electricity sales go to industrial customers. New interconnection requests, largely AI workloads, already exceed 500 MW. In parallel, the Pacific Gas and Electric Company (PG&E) reports a 10-GW data center interconnection pipeline as of August, up from 8.7 GW in May and 5.5 GW in February 2025.

The implications extend well beyond power system engineering. Data centers impose limited municipal service costs but generate significant property tax revenue and high-quality construction and operations jobs, often with strong labor standards. Their development catalyzes fiber deployment, substation and transmission upgrades, and electrification infrastructure that other customers can leverage.

However, lengthy permitting processes, local concerns, and grid bottlenecks are putting California's ability to keep and attract data center investment at risk. Neighboring states, such as Arizona, Nevada, and Utah, have been aggressively courting development by offering clearer timelines and lower delivered energy costs, as shown in Figure 1. If this infrastructure capital investment shifts elsewhere, the associated AI research and workforce clusters may migrate as well, weakening California's leadership in the AI economy.

The Electricity Trilemma in California

California's grid sits at the intersection of three competing imperatives: reliability, affordability, and sustainability, each sharpened by recent trends. Understanding how these pillars interact is essential to designing solutions that scale.

California's Clean Energy Leadership

California Senate Bill 100 sets California's clean energy transition goals: 60% renewables by 2030 and 100% clean electricity by 2045. By 2024, renewables supplied roughly 57% of in-state generation (solar 33%, wind 12%, hydro and geothermal 12%). Midday solar output now exceeds 20 GW, often pushing wholesale prices toward zero. During the spring months, CAISO can operate on 100% clean power for about 4–6 h per day. To integrate the variable generations, battery energy storage has scaled at record pace, from 0.5 GW in 2020 to 17 GW by November 2025. Storage now routinely delivers more than 6 GW during the evening peak, mitigating the "duck-curve" ramp and displacing fossil peakers. CAISO projects that battery capacity could exceed 20 GW by 2026. These milestones demonstrate a deeply transformed generation supply mix and system operation paradigm.

Reliability Under Climate Stress

California has experienced increasingly frequent and severe climate events, including extreme heat, prolonged drought, and widespread wildfires. These challenges have fundamentally reshaped how utilities plan and operate the electric system. To reduce wildfire risks, utilities now implement Public Safety Power Shutoffs, which can affect hundreds of thousands to millions of customers during peak fire seasons. As a result, overall reliability performance has declined over the past five years. At the same time, utilities are making major investments to harden the infrastructure against climate threats and enhance system resilience. Southern California Edison's proposed US$5.2 billion in wildfire mitigation spending from 2025 to 2028 highlights the scale and urgency of these adaptation efforts.

Another key reliability challenge is ensuring resource adequacy (RA) as renewable energy makes up a growing share of the generation mix. The August 2020 extreme heat wave revealed the vulnerability of the state's grid when CAISO directed rolling blackouts to maintain system stability. In response, the California Public Utilities Commission (CPUC) created the Slice of Day program, which requires load-serving entities to demonstrate enough resources to meet monthly peaks plus a reserve margin across all hours of the highest demand day. This program has driven billions of dollars of new investment, primarily in battery storage, which has significantly improved reliability, although at a high cost to ratepayers.

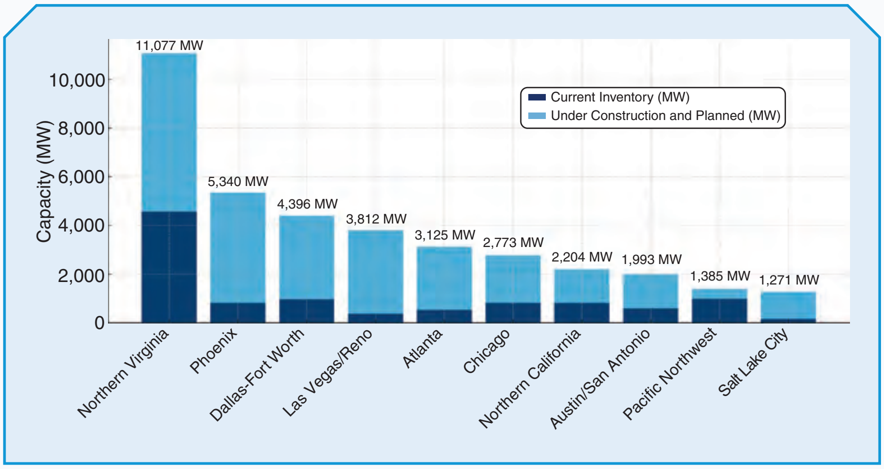

figure 1. The top-10 U.S. data center markets by future power capacity (current inventory plus under construction and planned). (Data source: Upwind and JLL 2024 Data Center Report.)

Affordability: The Pivotal Pillar

California's electricity costs rank among the highest in the nation. As of July 2025, average retail rates were approximately 32.58 ¢/kWh (residential), 29.94 ¢/kWh (commercial), and 25.31 ¢/kWh (industrial), compared with U.S. averages of 17.47 ¢/kWh, 14.15 ¢/kWh, and 9.29 ¢/kWh. In 2024, customers of the three major investor-owned utilities (IOUs) paid roughly US$54 billion for electric service, with revenue requirements projected to rise by 5%–9%

> AI demand may expose the limits of coordination and control within an already-complex system rather than create wholly new reliability phenomena.

in 2025. Two drivers loom large: wildfire mitigation and cost shifting associated with rooftop solar programs. Between 2019 and 2024, utilities collected roughly US$24 billion for wildfire mitigation and insurance premiums. Meanwhile, cost shifts associated with rooftop solar programs [net energy metering (NEM)/net billing tariff (NBT)] are estimated at about US$8.5 billion in 2024 alone, equivalent to 21%–27% of a typical nonparticipant residential bill. For 24/7 industrial customers, these figures raise competitiveness concerns and magnify the need for careful cost allocation and equitable benefit distribution. See Figure 2.

Do AI Energy Needs Intensify the Trilemma?

Whether AI fundamentally worsens the trilemma depends less on aggregate load volume than on where, when, and how demand materializes. Data centers concentrate large, steady loads; they can amplify local constraints or, if coordinated, improve utilization and lower average costs.

Reliability: Clustering and Disturbance Behavior

AI data centers often run at 90%–100% load factor and tend to cluster in a few corridors. This can strain transmission and distribution assets not designed for large blocks of continuous demand. CAISO's 2024–2025 transmission plan included US$4.8 billion in reliability-driven transmission projects and highlighted how rising electricity demand, especially in the Greater Bay Area, is driving the need for major transmission reinforcements. Operationally, the grid must account for how large facilities behave during disturbances. In a 2024 Northern Virginia event, roughly 1,500 MW of data center load tripped offline following a 230-kV equipment fault—raising questions about ride-through, controls, and coordination. California could see similar sensitivities as clusters grow, though modern interconnection standards and grid-code updates can mitigate risk. In this sense, AI demand may expose the limits of coordination and control within an already-complex system rather than create wholly new reliability phenomena.

Sustainability: Ambition Versus Round-the-Clock Reality

Many leading AI companies have made public climate commitments, and some have adopted a more ambitious 24/7 Carbon-Free Energy (CFE) target. Under the 24/7 CFE target, every kilowatt-hour of electricity consumed is matched with carbon-free electricity on an hourly basis, every day of the year. California's grid aligns well with the public climate commitments of major AI companies, but real-time 24/7 matching remains difficult. Constant demand can exceed renewable availability at night or during seasonal lulls, necessitating gas-fired balancing. Backup power complicates the picture: most large data centers still depend on diesel emergency generators for $N + 1$ reliability, even though air-quality rules sharply limit runtime. While the total energy involved is small, local impacts can be meaningful. Transition pathways—hydrogen-ready turbines, fuel cells, and long-duration storage—are promising but not yet universal.

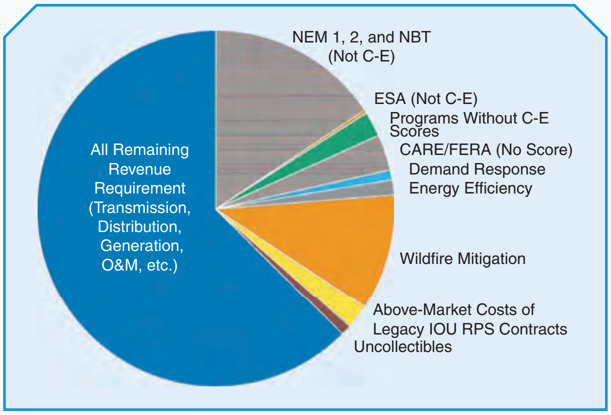

figure 2. The contribution to rates of programs funded through electric IOU revenues. NEM: net energy metering; C-E: cost-effectiveness; O&M: operation and maintenance; ESA: energy savings assistance; CARE: California Alternate Rates for Energy; FERA: Family Electric Rate Assistance; RPS: Renewables Portfolio Standard.

Affordability: Who Pays and When

Meeting new load requires substations and transmission expansions that can sum to tens of billions of dollars statewide. CAISO estimates about US$4.8 billion in near-term transmission projects; PG&E anticipates around US$73 billion in grid investments through 2030, reflecting rising demand across sectors rather than AI alone. Traditional regulatory frameworks spread these costs across all ratepayers. If AI load arrives faster than expected, the timing of cost recovery can create short-term bill pressure before longer term asset-utilization benefits appear. Conversely, well-sited, concentrated demand can improve utilization and lower average costs over time. The outcome is therefore uncertain—it is ultimately a function of siting, timing, and financing structures still in flux.

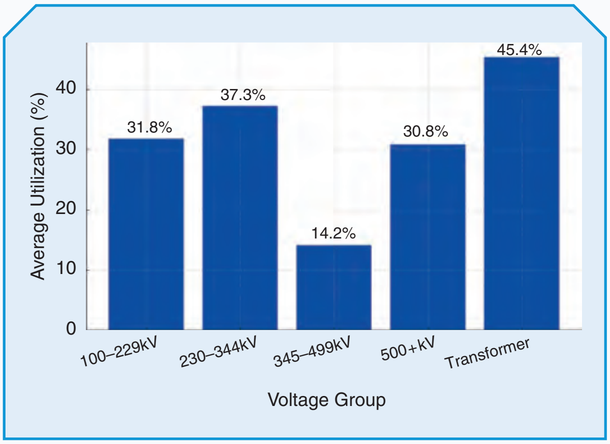

figure 3. California transmission utilization by voltage level with average load by voltage group under the 2025 summer-peak forecast with *N* – 1 contingencies. Many high-voltage lines are substantially underutilized, reflecting required reliability margins. (Analysis conducted by Stanford, based on WECC data.) WECC: Western Electricity Coordinating Council.

The bottom line: AI demand does not neatly worsen or ease the trilemma; it reconfigures it. Reliability risks hinge on clustering and controls; sustainability on hour-by-hour matching; affordability on cost allocation and timing. The next decade of choices will determine whether this evolution strains or strengthens system equilibrium.

From Trilemma to Triumph: A Technical Road Map

Turning AI data centers into a grid modernization catalyst requires three strategies: 1) unlock underutilized transmission capacity to power AI-driven load growth, 2) target clean and dispatchable resources to areas of highest grid impact, and 3) implement equitable cost-recovery mechanisms so that data centers fund needed upgrades without burdening other customers.

Increasing Transmission Utilization to Power AI Data Centers

Increase Transmission Utilization Spatially and Systemically

California's transmission grid presents a paradox. Despite substantial investments to integrate renewable generation, large portions of the network remain underutilized. Figure 3 shows the average utilization by voltage group across the state. This is not inefficiency but intentional design. Operators and planners maintain reliability margins to ensure the grid remains stable under contingency (*N* – 1) conditions. When a major line or transformer is out of service, power reroutes instantly, sometimes doubling the loading on adjacent elements, so the system must preserve headroom to stay secure during rare but critical events.

Under *N* – 1 contingency, average utilizations are just 31.8% for 115-kV lines, 37.3% for 230-kV lines, 14.2% for 345-kV lines, and 30.8% for 500-kV lines. Most high-voltage lines therefore operate well below their thermal ratings, with substantial apparent headroom across all voltage levels. Across the 1,145 230-kV lines in the state, ordering lines by maximum observed loading reveals a long, relatively flat stretch of low-to-moderate utilization (maximum utilization across full contingency ensemble) and only a small fraction of lines in a highly loaded tail, as shown in Figure 4. However, this headroom should

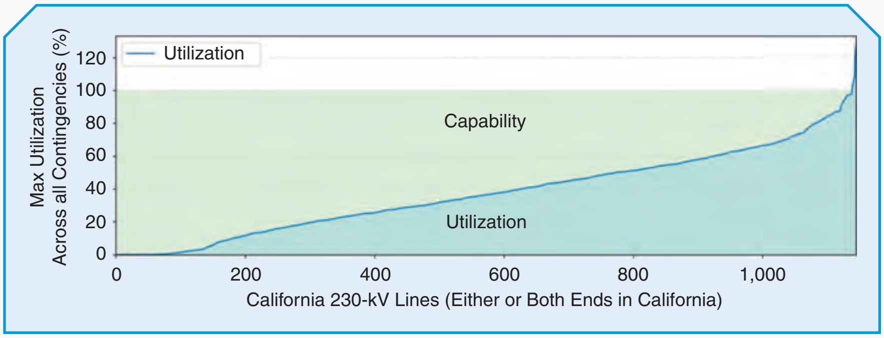

figure 4. Sorted utilization results for 230-kV transmission lines in CAISO.

not be directly interpreted as fully transferable thermal capacity. Voltage support limits, reactive power constraints, protection settings, and remedial action schemes can bind before conductor temperature does. Targeting this small subset of heavily loaded lines and their associated constraints is therefore essential for any strategy that aims to materially increase utilization across the broader network. Unlocking headroom safely requires stability-aware operating limits, and, in some corridors, enabling technologies, such as dynamic line ratings (DLRs), advanced power-flow controls including the flexible ac transmission system (FACTs), topology optimization, or other grid-enhancing technologies, are some of the fastest and most cost-effective ways to improve utilization and accommodate new large loads, such as data centers.

Recommendations are as follows:

- *Measure and publish planning-relevant utilization metrics*: Report annual $N-1$ loadings at important corridors, consistent with North American Electric Reliability Corporation Transmission Planning standards. Use these metrics to identify underused assets, binding constraints, and cost-effective upgrade candidates (such as grid-enhancing technologies and energy storage).
- *Develop load and headroom heatmaps*: Provide developers with sufficiently granular maps that show where capacity is available and where constraints bind to guide siting while respecting reliability. Steer interconnections toward headroom and away from congested points of interconnection (POIs).
- *Target upgrades to key constraints*: Focus infrastructure capital on the few constraints that limit many flows. Strategic upgrades including FACTs, DLRs, etc., can unlock substantial transfer capability, improve system-wide utilization, and interconnect new loads at a relatively modest cost.

Unlock Temporal Utilization

Typical transmission planning focuses on peak demand, which typically occurs for only about 1% of the year, as shown in Figure 5. Enhancing temporal flexibility is crucial for unlocking the underutilized grid capacity, allowing for better use of the current infrastructure without the immediate need for new power plants or significant infrastructure upgrades. Utilities are investing in flexibility with new technologies, such as energy storage and demand response programs. However, there's a big flexibility opportunity within data centers. Data centers have various dispatchable resources, like backup generators, uninterruptible power supply (UPS) batteries, and stationary energy storage. These facilities often keep these resources as backups, but they sit idle most of the time. Utilities can leverage these underutilized resources for grid flexibility, and this approach aligns closely with utilities' use of natural gas plants during peak demands. This shifts flexibility provision from expensive, seldom-used fossil fuel resources to more efficient, integrated data center resources.

Recent state legislation has begun promoting load shifting and dynamic pricing to improve the utilization of the existing grid infrastructure. The CEC has set a goal of achieving 7 GW of demand flexibility, encompassing load shifting and flexible consumption, by 2030. California's IOUs have launched pilot programs for distribution-level flexible interconnection, but progress so far has been limited. But, under a business-as-usual scenario, California is unlikely to reach the full 7-GW target without additional near-term policy and market interventions.

In today's overheated demand for AI computing power, data centers have little financial reason to use their facilities in demand response or flexibility programs. If utilities rewarded proven flexibility with faster interconnection and priority access to capacity, data centers' business goals would better align with the grid's needs, and

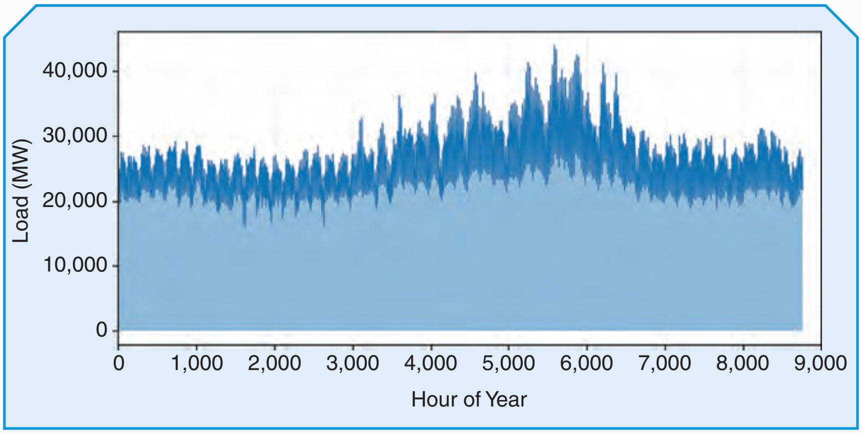

figure 5. CAISO 2025 8,760-h (12-month) hourly load profiles highlight the ~1% of hours that set annual peaks and the 99% with spare capacity—underscoring the value of temporal utilization. (Source: CPUC.)

> Awarding RA credit to compliant clean backup resources can help align private resilience investment with the system's need for dependable capacity.

supporting the grid would become a core part of how they operate, not just a side benefit.

Because AI data centers typically interconnect at the transmission level (100 kV and above), policy makers should expand flexibility incentives and mechanisms to cover the higher voltage. By turning idle data center dispatchable resources, such as backup power and storage systems, into active, grid-supportive flexibility, California can accelerate interconnection timelines, strengthen grid resilience, lower costs for all ratepayers, and advance its 7-GW flexibility goal.

Recommendations are as follows:

- *Transmission-level flexible interconnection*: California's IOUs have launched pilot programs for distribution-level flexible interconnections. We encourage regulators and utilities to extend these programs to the transmission level, enabling large flexible loads to connect quickly with clear operating frameworks.
- *Develop programs to utilize data center fast-ramping resources*: Data centers construct on-site backup power for reliability purposes, though they are rarely used (mostly for required annual testing purposes). A cost-effective program that promotes developing cleaner natural gas on-site units and batteries and that utilizes these resources during certain grid/operational conditions can untap an important resource that supports grid operations and provides value to all customers served by the grid. More importantly, data centers that are willing to provide grid flexibility should be rewarded with prioritized interconnection.
- *Standards, automation, and integration*: New technologies and data center flexibility management systems are needed to better integrate and coordinate behind-the-meter resources in response to grid needs. This includes deploying software solutions for real-time control, establishing standards for interoperability, and creating predictive algorithms for load shaping to better respond to grid conditions.

These measures convert idle capacity into reliable flexibility, advance California's 7-GW demand-flexibility procurement goals, accelerate interconnections, and reduce system costs for all customers.

Deploy Clean and Firm Capacity Where It Matters

California has rapidly expanded renewable energy and storage, making progress toward its goal of 100% clean electricity by 2045. However, the growth of data centers adds significant, continuous electricity demand. To sustainably meet this growing demand, California should accelerate the deployment of clean firm generation, such as geothermal energy, low-carbon fuel cells, and natural gas with carbon capture. Research indicates that combining renewables with clean firm generation substantially lowers overall system costs, land requirements, and transmission infrastructure needs compared with renewables-only scenarios. However, geothermal, low-carbon fuel cells, and carbon capture technologies remain early in commercialization and face stubbornly high costs. They also encounter prolonged permitting processes, financing hurdles, and community approval challenges that may delay projects beyond the timelines required by data centers. RA credits can improve bankability, but in most cases they will not by themselves close the cost gap with conventional diesel backup. Additional long-term contracts, targeted incentives, grants, tax credits, or other policy support will often be needed during early deployment.

Steering Load to Capacity and Deliverability

CAISO already had 185 GW in the queue before Cluster 15. Even after interconnection process improvements in early 2025, 118 GW of new generation capacity was resubmitted into the CAISO Cluster 15 queues, often without a clear path to deliverability. Figure 6 highlights in red areas where deliverability is scarce, yet queues are dense locations not typically aligned with data center siting. Strategic colocation of generation with load can accelerate both, particularly if the load interconnects at a POI that lacks generation deliverability. Coordinating load and generation interconnections can free up deliverability elsewhere, speeding both sides of the portfolio.

Recommendations are as follows:

- *Coordinated POI "express lanes"*: Launch a joint utility–CAISO program to evaluate paired interconnections, giving qualifying data centers with colocated clean firm generation a streamlined path: predefined operating envelopes for the load and expedited deliverability for the generator, plus fast-tracked siting, permitting, and interconnection at suitable sites with local engagement.
- *Market and grid integration*: Allow colocated clean firm resources to participate fully as front-of-meter resources (energy and ancillary services), using stacked revenues to improve project economics and drive adoption.
- *Better planning inputs*: Ensure that state resource and transmission planning explicitly reflect the scale, timing, and geography of emerging data center loads so that portfolios anticipate these needs rather than chasing them after the fact.

Strengthen Local RA

CAISO's 2026 Local Capacity Requirement study shows that the San Jose area requires 1,813 MW of capacity under the worst contingency, with a local deficiency of 943 MW. Figure 7 shows that deliverability remains available for new generation in San Jose. But, the region's dense, urban character, land-use constraints, permitting timelines, and likely community opposition may limit the siting of large-scale solar or wind projects. At the same time, expanding data center demand is driving higher urban load factors and increasing the need for resources that can produce reliably across many hours of the year. In this context, clean firm resources that are sited close to load will be essential to meeting local reliability needs, reducing dependence on fossil peakers, and supporting continued growth in electrification and AI-driven demand.

Recommendations are as follows:

- *Prioritize high-density, low-footprint, clean firm technologies in dense urban areas*: Establish procurement mechanisms that explicitly value local RA, with performance

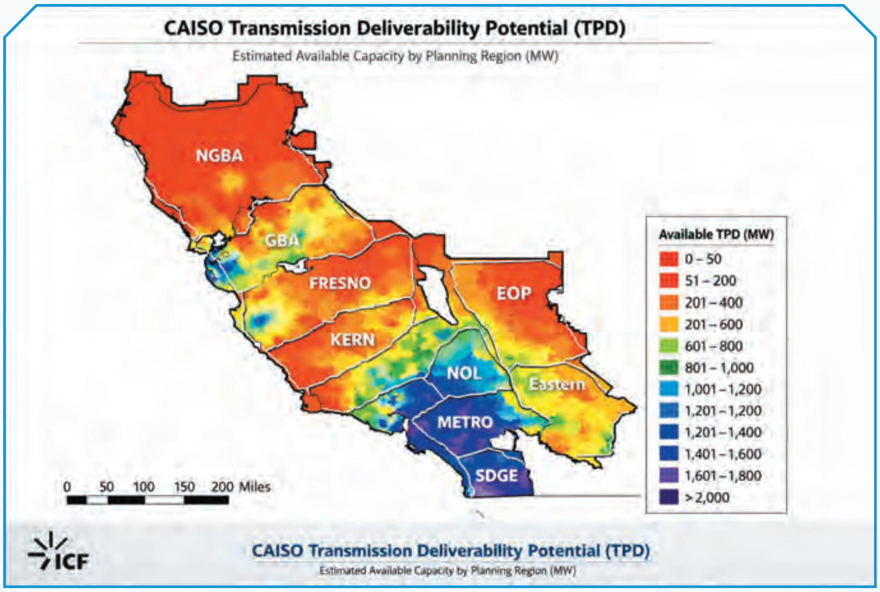

figure 6. CAISO deliverability status: red indicates areas lacking deliverability yet crowded with queued projects—poor matches for typical data center siting absent targeted coordination at POIs. (Source: ICF.)

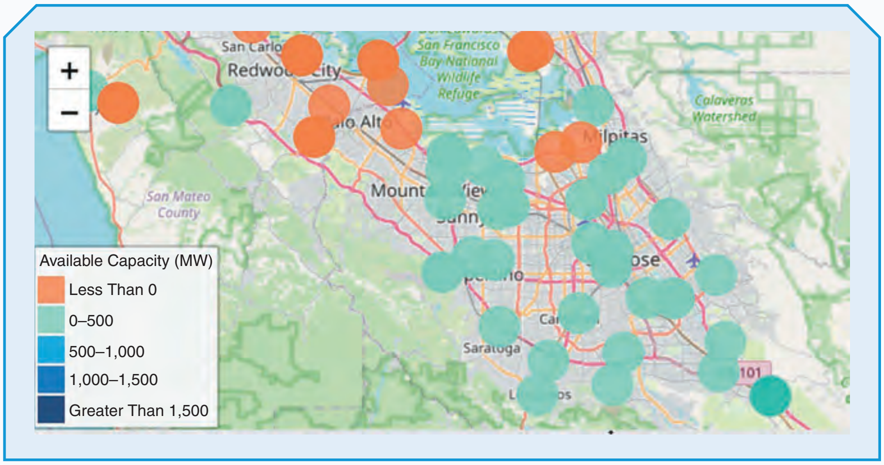

figure 7. San Jose deliverability and local RA context: sufficient deliverability remains for new generation.

> Utilities and regulators can require long-term capacity or minimum-payment commitments from large data centers, so that they pay for a defined level of service regardless of actual usage.

credits for fast-start/fast-ramp capability and disturbance support.

- *Allow compliant clean backup systems to earn local RA credit* where they meet emissions, availability, and dispatch standards—converting passive reliability insurance into active capacity that benefits all customers.

As Figures 3 and 4 illustrate, California's system has headroom, but it is unevenly distributed. Heatmaps and coordinated POI planning can guide large loads toward transmission capacity and free deliverability for generators. This reduces upgrade needs, shortens timelines, and lowers costs across the portfolio. This is particularly important where queues are long and equipment lead times are extended. In high-density nodes, such as San Jose, local RA needs are acute while land is scarce. The solution stack therefore leans toward clean firm technologies, advanced storage, and grid-supportive interconnection. Awarding RA credit to compliant clean backup resources can help align private resilience investment with the system's need for dependable capacity.

Ensure Data Centers Fund Upgrades Without Burdening Other Ratepayers

Rising data center demand will require substantial new generation and transmission investments. Many of these are projects that would not otherwise be built and that carry substantial fixed costs for utilities. Another risk is that, if demand projections fall short or major data centers close, utilities could be left with stranded assets that existing ratepayers must cover.

To manage this risk, utilities and regulators can require long-term capacity or minimum-payment commitments from large data centers, so that they pay for a defined level of service regardless of actual usage. These contracts, often 15–20 years in length, provide revenue certainty, support upfront financing of grid expansion, lower borrowing costs, and help ensure that AI-driven load growth supports overall affordability and reliability.

PG&E's interim Rule 30, approved by the CPUC in July, is a step in this direction because it clarifies how large, discrete loads will fund distribution-level upgrades. However, Rule 30 does not address the CAISO tariff transmission infrastructure. PG&E's 2025 Transmission Planning Process filing to CAISO highlights the scale of needed transmission-level upgrades and the need for complementary cost-recovery mechanisms that go beyond Rule 30.

Align Cost, Risk, and Timing: Affordability by Design

New demand can play a vital role in affordably modernizing California's energy infrastructure. By increasing total electricity consumption (in megawatt-hours), new large loads help distribute fixed system costs, such as transmission and distribution investments, wildfire mitigation expenses, and public purpose program surcharges, across more kilowatt-hours. This broader cost base can allow utilities to maintain, or even lower, average electricity rates for all customers. California faces an aging infrastructure challenge, with only a small portion of transmission lines operating near full capacity, while others remain significantly underutilized. Even without the data center boom, the state must still invest in upgrading and hardening its grid to meet safety and reliability standards. Under the current rate structure, these are investments that would likely drive rates upward. By introducing new, steady load growth, the state can modernize its infrastructure more affordably, leveraging economies of scale without increasing rates for existing customers.

That said, it is critical to distinguish between upgrades that are broadly needed for system modernization and those triggered specifically by new large-load interconnections, such as data centers. Hyperscalers and data center developers are willing to pay their fair share toward infrastructure investments required to serve them while improving overall grid utilization. Still, local communities remain cautious. If the AI-driven demand surge proves short-lived, they fear being left with stranded assets and long-term rate impacts. Establishing cost-recovery mechanisms, such as *take-or-pay* agreements or targeted infrastructure cost-sharing frameworks, can protect existing ratepayers by ensuring that those who drive new demand also shoulder the associated long-term costs.

Given their stringent reliability requirements, data centers typically install redundant on-site backup generation, most of which today relies on fossil fuel resources, such as diesel gensets. If these systems were replaced with or supplemented by clean and dispatchable generation, for example, hydrogen-fueled turbines, fuel cells, or long-duration storage, they could provide valuable local RA capacity. By allowing such clean backup resources to earn RA credit, California could enable utilities and community choice aggregators to procure local capacity without additional cost, while incentivizing data centers to invest in clean, dispatchable reliability assets. This approach would simultaneously advance the state's reliability and decarbonization goals, transforming backup generation from a passive reliability measure into an active contributor to a more resilient and sustainable grid.

What Success Looks Like

A successful path leverages AI demand to improve reliability, affordability, and sustainability simultaneously.

- *Reliability*: Large loads connect faster via transmission level flexible interconnection, operate within defined envelopes, and contribute fast-ramping flexibility through UPS batteries and on-site storage. Clean backup earns RA credit, and ride-through and ramp standards prevent sudden, destabilizing load loss during disturbances.
- *Affordability*: Fixed costs are spread across more megawatt-hours as new load comes online, targeted financing protects ratepayers from project-specific upgrades, transmission headroom is exploited before expensive expansions, and colocation with coordinated POIs avoids redundant build-outs.
- *Sustainability*: Clean firm resources deploy at priority POIs, diesel backup transitions toward lower emission alternatives, and workload orchestration aligns non-latency-sensitive compute with renewable-rich hours.

In this scenario, the system becomes 1) cleaner by harnessing flexible demand to integrate a higher penetration of renewables, 2) more affordable by using existing assets better and allocating costs fairly, and 3) more resilient by turning passive backup and idle headroom into active grid support. These outcomes are mutually reinforcing when planning, interconnection, and operations are aligned around measurable performance and transparent incentives.

Conclusion

California stands at a critical inflection point. AI-driven electricity demand is arriving faster than traditional infrastructure processes can accommodate. But that challenge is also an opportunity: by quantifying and utilizing transmission headroom, adopting flexible interconnections, and accelerating clean firm deployment at strategic nodes, the state can lower bills, improve reliability, and advance decarbonization while securing a competitive position in the global AI economy. If managed poorly or too slowly, demand will move elsewhere, and costs per unit will rise for the customers who remain. A practical path forward treats large new loads as partners in meeting statewide goals, not as problems to be kept at a distance.

For Future Reading

"Powering AI at speed in California," Stanford Precourt Inst. for Energy, Stanford, CA, USA, Sep. 2025. [Online]. Available: https://drive.google.com/file/d/14NhXPKynQCjbBC99xVk6tA5pqCD9U052/view

"CPUC response to executive order N-5-24," California Public Utility Commission, San Francisco, CA, USA, Feb. 2025. [Online]. Available: https://www.cpuc.ca.gov/-/media/cpuc-website/industries-and-topics/reports/cpuc-response-to-executive-order-n-5-24.pdf

J. C. S. Long et al., "Clean firm power is the key to California's carbon-free energy future," Issues in Sci. and Technol., Tempe, AZ, USA, Mar. 2021. [Online]. Available: https://issues.org/california-decarbonizing-power-wind-solar-nuclear-gas/