American Journal of Student Research

Original Research Article

AI Data Center Expansion and Its Implications for Energy Demand and Environmental Justice: Evidence from New Jersey and Comparative U.S. Regions

Katie Chung

Bergen County Academies, 200 Hackensack Ave, Hackensack, NJ 07601, United States

ABSTRACT

The rapid growth of Artificial Intelligence (AI) has sharply increased demand for data center infrastructure, raising concerns about energy use, greenhouse gas emissions, and environmental justice. This study examines AI-driven data center development in New Jersey, estimating effects on electricity demand, air pollution, and public health using a transparent screening framework. Facility-specific capacity data are compiled from publicly accessible sources, EPA eGRID 2022 average emission rates are applied under two utilization scenarios, public health effects are assessed using the EPA COBRA screening model, and environmental justice implications are evaluated using EPA EJScreen and NJDEP overburdened-community mapping tools. Screening results indicate that New Jersey's disclosed data center capacity could expand 4.16-fold, increasing estimated electricity demand from approximately 2.06–2.74 TWh per year (2.8–3.7% of statewide retail sales) to 8.55–11.40 TWh per year (11.6–15.5% of statewide retail sales). Associated CO_2 emissions under full build-out are estimated at approximately 2.55–3.40 million metric tons per year, with modeled health damages reaching approximately \$271–\$450 million annually under high utilization. Several existing and planned facilities are located in or near communities already identified as socioeconomically vulnerable or facing elevated pollution burdens. Comparative analysis with Northern Virginia and Central Ohio illustrates how rapid, large-scale data center growth creates grid reliability, cost allocation, and environmental equity challenges in the absence of proactive policy frameworks. These findings support improved facility-level disclosure, targeted tariff design, cleaner incremental electricity supply, and rigorous tract-level environmental justice review.

Keywords: Artificial Intelligence; Data Centers; Energy Demand; Environmental Justice; Air Pollution; Public Health; Infrastructure

INTRODUCTION

The rapid expansion of artificial intelligence (AI) has sharply increased the demand for data processing and storage infrastructure. Training and operating large AI models now requires substantial amounts of electricity. The International Energy Agency estimates that data centers consumed approximately 415 TWh globally in 2024, roughly 1.5% of world electricity use, and projects that global data center demand could rise to approximately 945 TWh by 2030 in its base scenario. In the United States, Lawrence Berkeley National Laboratory estimated that data centers consumed approximately 176 TWh in 2023, equal to about 4.4%

Corresponding author: Katie Chung, E-mail: kjchung0803@gmail.com.

Accepted June 3, 2026
https://doi.org/10.70251/HYJR2348.43506517

of total U.S. electricity consumption, with projections ranging from roughly 325 to 580 TWh by 2028.

These trends carry particular significance in states such as New Jersey, where data center growth is occurring within a densely populated, heavily interconnected electricity system and where environmental permitting and environmental justice are already salient policy concerns. New Jersey's Environmental Justice Law requires the Department of Environmental Protection (DEP) to evaluate environmental and public-health impacts for certain covered facilities sited in overburdened communities, and the state has considered large-load tariff legislation intended to protect other utility customers from costs associated with exceptionally large new electricity users. According to the U.S. Energy Information Administration, statewide retail electricity sales total approximately 73.53 million MWh annually, providing a concrete baseline against which incremental data center demand can be assessed.

New Jersey is emerging as a location for large-scale AI data center development. The state already hosts numerous facilities in areas such as Secaucus, Newark, and Piscataway, and new hyperscale projects, including a planned 300-MW facility in Vineland, are under development. These developments raise questions about the adequacy of New Jersey's energy grid, water resources, and environmental protections given a rapid increase in projected demand.

This study asks three focused questions. First, how much annual electricity demand could the currently installed and publicly announced major New Jersey data center inventory represent under transparent operating assumptions? Second, what average-grid emissions and screening-level public-health burdens are implied by those demand scenarios? Third, do the facilities and expansion areas identified in the inventory intersect with overburdened communities or elevated environmental justice indicators in ways that merit closer review? This paper is intentionally framed as a scenario-based screening analysis of major disclosed facilities rather than a definitive causal or statewide census study, in keeping with the limitations of the publicly available data.

LITERATURE REVIEW

Energy and Environmental Impact of AI Data Centers

Research shows that the shift from conventional cloud computing to AI-intensive workloads has significantly increased data center electricity demand. Training large-scale models requires extensive computation on GPUs and TPUs, consuming substantial amounts of energy. GPT-3 training has been estimated at approximately 1,287 MWh, while GPT-4 has been estimated to exceed 5,000 MWh. AI inference adds continuous demand, with some analyses suggesting that individual AI queries use substantially more electricity than conventional web searches.

Global data center electricity consumption is projected to rise sharply with continued AI growth. Data centers in major hubs such as Loudoun County, Virginia, already consumed approximately 21% of local electricity as of 2023, illustrating the local grid significance of geographically concentrated demand.

Cooling and water use represent additional environmental burdens. Typical data centers operate at Power Usage Effectiveness (PUE) values of 1.5 to 2.0, indicating substantial overhead energy beyond the IT load itself. Water use for evaporative cooling adds further environmental pressure, though precise facility-level water consumption data are rarely disclosed publicly. Because no facility-specific cooling-technology data are available for the New Jersey facilities examined in this study, water use is discussed here as an important dimension of data center environmental impact that is outside the present empirical boundary and is recommended as a priority for future work.

Beyond operational energy use, AI infrastructure also generates environmental impacts through hardware production and disposal. Construction embeds emissions in materials such as steel and concrete, while high-performance chips depend on resource-intensive supply chains and contribute to increasing levels of electronic waste.

Fossil-Fuel Reliance and Grid Carbon Intensity

Despite sustainability commitments, most AI data centers continue to operate on grids that rely substantially on fossil fuels. Because marginal electricity supply in many regions is generated from natural gas or coal, the emissions associated with incremental data center demand can be substantial. Companies frequently rely on renewable energy credits (RECs) and off-site power purchase agreements, which can enable "100% renewable" claims even when the facility draws electricity from fossil-heavy grids. A 2024 investigation estimated that actual emissions from major technology firms may be substantially higher than those reported in corporate sustainability disclosures. Backup diesel generators further contribute to local air pollution, particularly during outages, maintenance testing, or periods of peak grid demand.

Air Pollution and Public Health Outcomes

In addition to CO_2 emissions, data center energy use contributes to the release of $PM_{2.5}$, NO_x, SO_2, and other pollutants with significant public health consequences. A 2024 preprint by Han, Wu, Li, Wierman, and Ren introduced a methodology to model lifecycle pollutant emissions from data centers and computing tasks, estimating that by 2030, pollution associated with U.S. data center growth could be associated with approximately 1,300 premature deaths annually and over 600,000 respiratory symptom cases, with total modeled health damages approaching $20 billion per year. These impacts are modeled as unevenly distributed, with communities near power plants or large-scale data center developments potentially experiencing higher pollutant exposure, and with modeled effects extending across multiple states through $PM_{2.5}$ transport.

Environmental Justice and Siting Inequalities

Existing research indicates that data center development may follow broader patterns of environmental inequality. National spatial analyses have found that a substantial share of U.S. data centers are located in census tracts with above-median social vulnerability, suggesting that the environmental burdens associated with data center expansion may disproportionately affect already-stressed communities.

Procedural justice concerns also arise, as data centers are frequently categorized as “clean” infrastructure and may not undergo the same level of permitting scrutiny as traditional industrial developments. This can limit community input during siting and reduce transparency around potential environmental impacts.

Greenwashing, ESG Claims, and Transparency Gaps

Corporate sustainability narratives often obscure the localized environmental impacts of AI data centers. Firms frequently emphasize efficiency metrics such as PUE and highlight renewable energy procurement while omitting absolute increases in energy demand, site-level emissions, and water use. Renewable energy credits play a central role in enabling “100% renewable” claims despite discrepancies between reported and actual energy sources. Many operators do not publicly disclose facility-level emissions of NO_x or $PM_{2.5}$, nor detailed water consumption data. These transparency gaps have led to increasing calls for mandatory reporting standards, stronger permitting requirements, and policies that tie sustainability claims to verifiable, real-time environmental performance.

METHODS AND MATERIALS

Study Design and Regional Focus

This study uses a comparative mixed-methods approach to assess the environmental, health, and equity implications of AI-driven data center development. New Jersey serves as the primary case study, with Northern Virginia and Central Ohio used as benchmark regions. Northern Virginia represents a mature, high-density data center hub, while Central Ohio reflects a rapidly expanding hyperscale market. Within New Jersey, the analysis includes operational, under-construction, and planned data centers with publicly disclosed capacity information, with attention to clustering patterns across northern and southern regions. This design enables evaluation of how current development trends and policy frameworks may shape future environmental and equity outcomes. The facility inventory is explicitly defined as a major-facility, public-disclosure-based compilation of large sites suitable for scenario analysis of grid impacts, and should not be interpreted as a full census of all enterprise server rooms or undisclosed colocation sites.

Data Center Identification and Classification

Facility-level data for New Jersey were compiled using DatacenterMap as the primary source, supplemented by Baxtel and direct company disclosures. These sources were selected because no single official statewide registry of New Jersey data centers with validated capacity information currently exists; DatacenterMap and Baxtel represent the most comprehensive publicly accessible aggregators of facility-level capacity data, and their use is consistent with recent national research that likewise relied on public and proprietary market datasets. The inventory should therefore be interpreted as a representative major-facility compilation rather than a complete enumeration of all sites. Data centers were classified by operational status, operator type, and facility size. Both hyperscale facilities and large campus-style or carrier-neutral developments were included; smaller enterprise-only facilities were excluded unless part of larger campuses. Each site was geographically assigned to a county, which serves as the primary spatial unit of analysis for electricity demand, emissions, and health impact modeling.

For Northern Virginia and Central Ohio, site-level identification was not conducted; instead, regional statistics and published studies were used for comparative analysis.

Estimation of Electricity Demand

Annual electricity consumption was estimated from disclosed IT capacity using the standard screening expression: *Annual electricity consumption (MWh) = IT Load (MW) × Load Factor × 8,760 × PUE.* A PUE value of 1.5 was applied as a transparent screening midpoint. This value is consistent with recent U.S. aggregate estimates, which place average data center PUE near 1.4 in 2023, and with the Uptime Institute's reported industry average of 1.58 for the same year; 1.5 therefore sits between these two reference points. Two load factor scenarios were modeled to capture variation in sustained utilization: 0.6 (moderate utilization) and 0.8 (high utilization). These are scenario bounds, not probability-weighted forecasts.

Emissions Estimation Using EPA eGRID

Estimated electricity consumption was converted into emissions using EPA eGRID 2022 average output emission rates. Because New Jersey operates within the PJM interconnection, the RFCE (ReliabilityFirst Corporation East) subregion rates were applied for the primary scenario results, consistent with the facility inventory. It is important to note that eGRID reports average output emission rates, not marginal emission rates; the estimates reported here are therefore average-grid attributional screening estimates rather than marginal-cost or dispatch-based forecasts.

Emissions were calculated as: *Metric tons/yr = (MWh × emission factor [lb/MWh]) / 2,204.6*

Emission factors applied were: CO_2 at 657.4 lb/MWh, NO_x at 0.300 lb/MWh, and SO_2 at 0.302 lb/MWh (EPA eGRID 2022, RFCE subregion). All results are reported in metric tons per year, using the conversion formula above. A sensitivity check using New Jersey state-average eGRID 2022 output emission rates was also conducted to assess the sensitivity of results to the geographic averaging choice. Results of this sensitivity analysis are reported in Section 3.3.

Health Impact Assessment

Public health impacts were estimated using the EPA COBRA (Co-Benefits Risk Assessment) screening model, which links changes in emissions to estimated changes in $PM_{2.5}$ concentrations and associated health outcomes. COBRA outputs include estimates of premature mortality, asthma exacerbations, hospital admissions, and monetized health damage. Health impacts were modeled at the county level, consistent with COBRA's spatial resolution. All results are described explicitly as screening estimates from a reduced-form model rather than observed outcomes; the reported ranges reflect the electricity-demand scenarios combined with COBRA's own valuation ranges, not statistical confidence intervals.

Environmental Justice and Spatial Analysis

To assess equity implications, data center locations were compared with demographic and environmental indicators from EPA EJScreen Version 2.3 and the NJDEP overburdened-community mapping tool (EJMAP). Variables considered include income, race and ethnicity, cumulative pollution burden, and asthma prevalence where available. Because this study does not include a tract-by-tract dataset with matched control communities, the environmental justice analysis is explicitly descriptive and screening-oriented. A more rigorous extension would geocode every facility boundary, extract host-tract and adjacent-tract EJ variables and compare them with state medians or matched control tracts using nonparametric tests or logistic regression. That tract-level statistical analysis is recommended for future work but could not be completed from the materials available for this study.

Policy Context and Comparative Interpretation

A qualitative policy analysis was conducted to contextualize empirical findings. For New Jersey, this included review of the state's Environmental Justice Law, zoning practices, and permitting requirements for backup generators, as well as proposed state legislation. Federal policies, including incentives under the Inflation Reduction Act, were also considered. Evidence from Northern Virginia and Central Ohio, including documented reliability events, regulatory responses, and market conditions, was used to inform interpretation of New Jersey's development trajectory.

RESULTS

Identification and Capacity of Data Centers in New Jersey

A total of sixteen major data center facilities and campuses were identified across New Jersey using the facility inventory described in Section 2.2. Facilities span a range of operational statuses, from fully operational carrier hotels and multi-tenant campuses to large hyperscale developments currently under construction. Operational facilities include 165 Halsey Street (Newark), CoreSite NY3 (Secaucus), CyrusOne NYM1 (Somerset), QTS Piscataway, and several others concentrated in northern New Jersey. Large hyperscale developments under construction or in planning include the CoreWeave campus in Kenilworth (250 MW), Earth Station-76 in Gloucester County (150 MW), and the Nebius AI campus in Vineland, Cumberland County (up to 300 MW). Detailed facility-level data are provided in Table 1.

Across all operational sites, total installed IT capacity equals 260.9 MW. Under full build-out of planned and under-construction facilities, total statewide disclosed IT capacity rises to 1,084.4 MW, representing a 4.16-fold increase. The majority of new capacity is concentrated in three southern and central New Jersey counties: Cumberland, Union, and Gloucester.

Current and Projected Electricity Demand

Under current operational capacity, estimated electricity demand ranges from approximately 2.06 TWh per year (load factor 0.6) to 2.74 TWh per year (load factor 0.8). Relative to New Jersey's annual retail electricity sales of 73.53 million MWh, this corresponds to approximately 2.8–3.7% of statewide retail electricity consumption.

Under the full disclosed build-out scenario, estimated demand increases to approximately 8.55 TWh per year (load factor 0.6) to 11.40 TWh per year (load factor 0.8), equivalent to approximately 11.6–15.5% of statewide retail electricity sales. These figures represent scenario estimates based on disclosed capacity and stated modeling assumptions, not predictions of actual future consumption.

Emissions Implications of Data Center Electricity Demand

Using EPA eGRID 2022 RFCE average output emission rates, the currently installed inventory is estimated to correspond to approximately 0.61–0.82 million metric tons of CO_2 per year, approximately 254–339 metric tons of NO_x per year, and approximately 256–

Table 1. *Inventory of major data center facilities in New Jersey, including operational, expansion, and planned sites. Built-out IT capacity (MW), total potential capacity, and expansion runway are shown. Data compiled from DatacenterMap, Baxtel, and company disclosures. PUE values standardized at 1.5 unless otherwise specified.*

Facility	County	Status	Built-out IT MW	Total Potential IT MW	Runway MW	PUE
165 Halsey Street	Essex	Operational	23.5	82.0	58.5	1.5
365 Data Centers NJ3	Bergen	Operational	18.0	18.0	0	1.5
CoreSite NY3	Hudson	Operational	15.0	15.0	0	1.5
CyrusOne NYM1	Somerset	Operational	40.0	40.0	0	1.5
H5 Data Centers NJ	Hudson	Operational	5.0	5.0	0	1.5
Iron Mountain NJE-1	Middlesex	Operational	30.0	30.0	0	1.5
QTS East Windsor 1	Mercer	Operational	28.0	70.0	42.0	1.5
QTS Piscataway 1	Middlesex	Operational	52.0	67.0	15.0	1.5
NJFX (existing)	Monmouth	Operational	9.0	9.0	0	1.5
NJFX AI Hall Expansion	Monmouth	Expansion	0	8.0	8.0	1.25
Csquare EWR2	Hudson	Operational	20.7	20.7	0	1.5
Csquare EWR3	Middlesex	Operational	8.1	8.1	0	1.5
Csquare EWR5	Hudson	Operational	11.6	11.6	0	1.5
CoreWeave Kenilworth	Union	Under Construction	0	250.0	250.0	1.5
Earth Station-76 DC	Gloucester	Under Construction	0	150.0	150.0	1.5
Nebius NJ DC	Cumberland	Under Construction	0	300.0	300.0	1.5
TOTAL			**260.9**	**1,084.4**	**823.5**	—

341 metric tons of SO_2 per year.

Under the full disclosed build-out scenario, these estimates rise to approximately 2.55–3.40 million metric tons of CO_2 per year, approximately 1,056–1,408 metric tons of NO_x per year, and approximately 1,063–1,417 metric tons of SO_2 per year. County-level emissions inputs for health modeling are shown in Table 2 (load factor 0.6) and Table 3 (load factor 0.8). In both scenarios, Cumberland, Union, and Gloucester counties account for the largest incremental emissions, corresponding to the locations of the planned hyperscale campuses.

A sensitivity analysis using New Jersey state-average eGRID 2022 output emission rates yields lower estimates. Under this alternative, the currently installed inventory corresponds to approximately 0.45–0.61 million metric tons of CO_2 per year, 170–226 metric tons of NO_x per year, and 57–76 metric tons of SO_2 per year. Under the full build-out scenario, the corresponding estimates are approximately 1.89–2.52 million metric tons of CO_2 per year, 704–938 metric tons of NO_x per year, and 236–314 metric tons of SO_2 per year. Relative to the RFCE-based results, these state-average estimates are lower by approximately 26% for CO_2, 33% for NO_x, and 78% for SO_2. Full sensitivity results are summarized in Table 4. The RFCE-based results should be interpreted as one reasonable regional-average screening case rather than a unique forecast.

Geographic Distribution of Data Center Growth

Operational data centers are concentrated in northern New Jersey, particularly in Middlesex, Hudson, Somerset, and Essex counties, reflecting historical telecommunications infrastructure and proximity to New York City. This distribution is shown in Figure 1.

Future expansion is concentrated in central and southern regions, driven by large hyperscale developments. The projected full build-out distribution is illustrated in Figure 2, which shows a clear geographic

Table 2. *Estimated county-level increases in electricity consumption and associated emissions under the load factor 0.6 scenario. Emissions calculated using EPA eGRID 2022 RFCE average output emission rates.*

County	Δ MWh/yr	Δ CO_2 (metric tons/yr)	Δ NO_x (metric tons/yr)	Δ SO_2 (metric tons/yr)
Cumberland (Nebius)	2,365,200	705,290	321.9	324.0
Union (CoreWeave)	1,971,000	587,742	268.2	270.0
Gloucester (Earth Station-76)	1,182,600	352,645	160.9	162.0
Essex (165 Halsey expansion)	461,214	137,532	62.8	63.2
Mercer (QTS East Windsor)	331,128	98,741	45.1	45.4
Middlesex (QTS Piscataway)	118,260	35,265	16.1	16.2
Monmouth (NJFX AI Hall)	52,560	15,673	7.2	7.2
Bergen / Hudson / Somerset	0	0	0	0

Table 3. *Estimated county-level increases in electricity consumption and associated emissions under the load factor 0.8 scenario.*

County	Δ MWh/yr	Δ CO_2 (metric tons/yr)	Δ NO_x (metric tons/yr)	Δ SO_2 (metric tons/yr)
Cumberland (Nebius)	3,153,600	940,387	429.1	432.0
Union (CoreWeave)	2,628,000	783,656	357.6	360.0
Gloucester (Earth Station-76)	1,576,800	470,193	214.6	216.0
Essex (165 Halsey expansion)	614,952	183,375	83.7	84.2
Mercer (QTS East Windsor)	441,504	131,654	60.1	60.5
Middlesex (QTS Piscataway)	157,680	47,019	21.5	21.6
Monmouth (NJFX AI Hall)	70,080	20,897	9.5	9.6
Bergen / Hudson / Somerset	0	0	0	0

Table 4. *Sensitivity analysis using New Jersey state-average eGRID 2022 output emission rates. All values in metric tons per year (CO_2 in million metric tons/yr). Values shown for both currently installed and full build-out scenarios under load factors of 0.6 and 0.8.*

Scenario	Load Factor	CO_2 (M metric tons/yr)	NO_x (metric tons/yr)	SO_2 (metric tons/yr)	Note
Currently installed (260.9 MW)	0.6	0.45	170	57	NJ state-avg eGRID 2022
Currently installed (260.9 MW)	0.8	0.61	226	76	NJ state-avg eGRID 2022
Full build-out (1,084.4 MW)	0.6	1.89	704	236	NJ state-avg eGRID 2022
Full build-out (1,084.4 MW)	0.8	2.52	938	314	NJ state-avg eGRID 2022

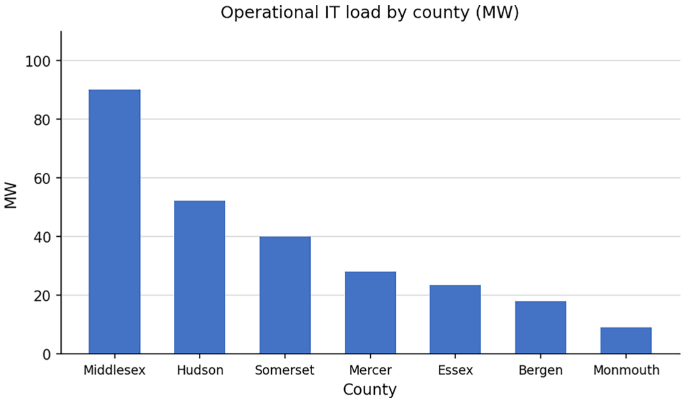

Figure 1. *Distribution of operational data center IT load (MW) across New Jersey counties. The concentration of existing facilities in northern counties reflects historical telecommunications infrastructure and proximity to New York City.*

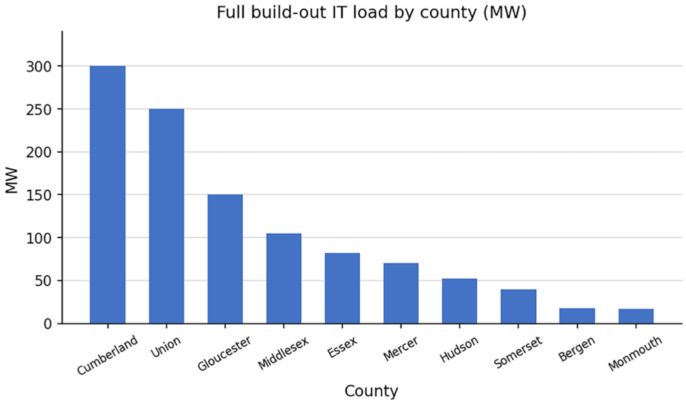

Figure 2. *Projected full build-out IT load (MW) by county. The figure illustrates a geographic shift toward southern and central New Jersey, driven by large hyperscale developments in Cumberland, Union, and Gloucester counties.*

shift toward Cumberland, Union, and Gloucester counties. These three counties account for approximately 700 MW — the majority of new disclosed capacity under the full build-out scenario.

Public Health Screening Results (COBRA)

Health impacts associated with projected emissions changes were estimated using the EPA COBRA screening model. Under the load factor 0.6 scenario, total estimated health damages associated with the expansion range from approximately \$205 million to \$339 million annually, with estimated premature mortality of approximately 13.0 to 20.2 deaths per year. Under the load factor 0.8 scenario, estimated health damages increase to approximately \$271 million to \$450 million annually, with premature mortality estimates rising to approximately 17.2 to 29.2 deaths per year. These ranges reflect COBRA's own valuation methodology combined with the demand scenario bounds and should be read as screening-level estimates rather than precise forecasts. County-level distributions of monetized health damages are shown in Figures 3 and 4. In both scenarios, the largest estimated damages are concentrated in Cumberland, Union, and Gloucester counties.

Environmental Justice Screening

Data center locations were compared with demographic and environmental indicators using EPA EJScreen and NJDEP's EJMAP tool. This comparison identifies two patterns at the screening level. First, many existing operational facilities are concentrated in northern New Jersey — particularly the Newark–Secaucus corridor across Essex, Hudson, and Middlesex counties — where many census tracts already experience elevated environmental burdens associated with transportation, industrial activity, and freight

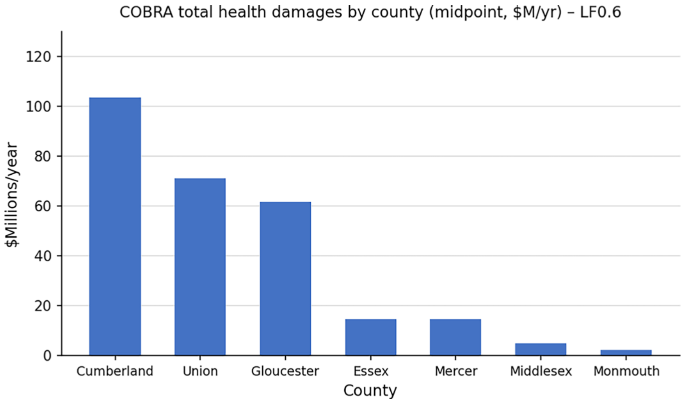

Figure 3. *Estimated annual monetized health damages by county under the load factor 0.6 scenario (midpoint values), based on EPA COBRA modeling. Damages reflect estimated impacts from increased $PM_{2.5}$ formation due to NO_x and SO_2 emissions.*

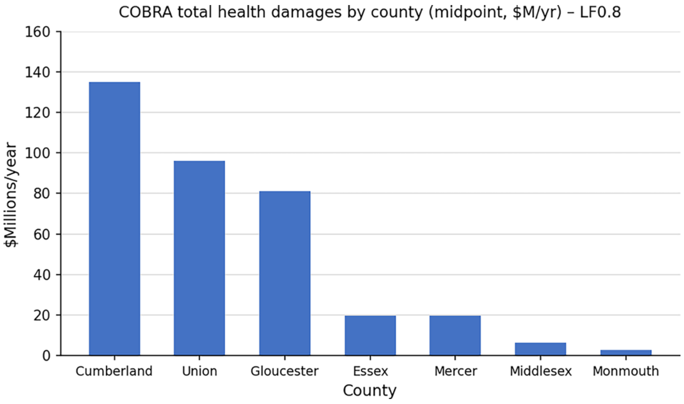

Figure 4. *Estimated annual monetized health damage by county under the load factor 0.8 scenario (midpoint values). Higher utilization results in larger modeled emissions and correspondingly larger estimated health burdens in counties hosting major expansions.*

infrastructure. Second, the largest planned hyperscale developments are concentrated in southern and central New Jersey, particularly in Cumberland and Gloucester counties, where some surrounding communities exhibit notable socioeconomic vulnerability according to EPA EJScreen indicators.

These findings are screening-level and do not establish a statistically demonstrated causal siting pattern. The stronger claim supported by this analysis is narrower: some of the growth areas identified in the inventory overlap with communities that already merit closer scrutiny under NJDEP overburdened-community screening and EPA EJScreen indicators, warranting more rigorous tract-level analysis.

DISCUSSION

Scale of New Jersey's Data Center Expansion

The screening results indicate that major data center expansion could become a consequential load-growth driver in New Jersey even when analyzed conservatively as a public-disclosure-based inventory rather than a full statewide census. The most policy-relevant finding is not any single emissions total, but the scale of the potential electricity demand increment itself. An estimated increase equivalent to approximately 11.6–15.5% of current statewide retail electricity sales under full disclosed build-out would be large enough to affect generation planning, transmission investment, rate design, and environmental permitting — consistent with broader national findings that data centers have become a meaningful contributor to electricity-demand growth.

Two distributional patterns in these results are particularly important for planning purposes. First, the magnitude of potential load growth is sufficient to reshape transmission planning and generation requirements, suggesting the need for integration into regional grid forecasting. Second, the geographic distribution of development is shifting from established northern New Jersey facilities to Cumberland, Union, and Gloucester counties — the same counties that account for the largest modeled health damages in the COBRA screening. This concentration of projected impacts suggests that environmental and public health burdens may be regionally concentrated rather than evenly distributed across the state.

Northern Virginia as a Mature Hyperscale Comparator

Northern Virginia provides a useful benchmark as the world's most mature data center market. The region hosts approximately 5,050 MW of data center electricity demand, largely concentrated in Loudoun County, and accounts for a substantial share of global data center capacity. The Virginia Joint Legislative Audit and Review Commission (JLARC) documented reliability risks at this level of density, including an incident in which approximately 60 facilities simultaneously disconnected, removing roughly 1,500 MW of load and requiring operator intervention to prevent broader system instability.

This precedent is relevant for New Jersey, where the facility inventory includes multiple planned campuses in the 150–300 MW range. Clustering of facilities at this scale within a shared transmission network increases the likelihood of synchronization or disturbance events. Northern Virginia also contextualizes the magnitude of modeled health damages: research by PSE Healthy Energy estimated that backup generator emissions alone in Northern Virginia may produce approximately $220–$300 million in annual health damages under moderate scenarios — a range similar in magnitude to the COBRA estimates reported here for New Jersey's expansion scenario, suggesting that total health-related costs could be larger if generator use increases under grid stress conditions.

Central Ohio as a Rapid-Growth Policy Case

Central Ohio serves as a comparator for policy adaptation in response to rapid hyperscale expansion. Regional capacity grew substantially from early 2024 through year-end, with projections indicating continued multi-gigawatt-scale growth. In response, the Public Utilities Commission of Ohio (PUCO) required development of data-center-specific tariff structures designed to protect existing ratepayers from infrastructure costs. These frameworks include minimum bill provisions, financial assurance requirements, and exit fees to ensure that large customers bear the financial risks associated with capacity commitments.

This approach is directly relevant for New Jersey. The projected expansion implies substantial transmission investment and potential generation expansion. Without targeted tariff structures, these costs could be shifted to residential and commercial customers. Ohio's experience also highlights the risks of speculative demand: contracts for data center load have exceeded operational capacity in some cases, creating uncertainty for infrastructure planning if projects are delayed or canceled — underscoring the importance of aligning private development incentives with public investment risk.

New Jersey's Data Center Trajectory

The cross-case comparisons reframe New Jersey's emerging data center landscape in several ways. First, the estimated potential demand share of approximately 11.6–15.5% of statewide electricity consumption indicates that data centers may need to be treated as a major energy planning sector, not a marginal one. Both Northern Virginia and Central Ohio demonstrate that once demand reaches this threshold, transmission planning, reliability management, and cost allocation become central policy challenges.

Second, the spatial distribution of New Jersey's development suggests that environmental and public health burdens could be concentrated in specific counties. County-level emissions modeling concentrates impacts in Cumberland, Union, and Gloucester counties. Unlike Northern Virginia's highly concentrated single-cluster model, New Jersey's development pattern spans multiple regions, potentially broadening distributional impacts across a more geographically dispersed area.

Third, the modeled increases in NO_x and SO_2 contribute to secondary $PM_{2.5}$ formation, producing near-term health impacts alongside longer-term carbon concerns. These findings align with emerging research emphasizing the localized air quality implications of large electricity loads and associated backup generator infrastructure.

Policy Implications for New Jersey

The findings suggest that policy responses should address three interrelated challenges: grid infrastructure costs, emissions and health impacts, and environmental justice.

On grid cost allocation and infrastructure planning: New Jersey policymakers have already begun considering regulatory frameworks for large electricity loads. Proposed legislation such as Senate Bill S731 would require electric utilities to establish special tariff structures for certain large-load customers, including data centers, to ensure that infrastructure costs are not shifted to other ratepayers. The screening results presented here support the rationale for such policies; effective frameworks should include minimum bill provisions, financial assurance requirements, and exit fees consistent with lessons from Ohio.

On emissions and generator governance: The Virginia experience demonstrates that generator use can significantly increase local air pollution impacts. Given the estimated health damages in the $205–$450 million range under the expansion scenario, New Jersey regulators may need to treat emissions mitigation as a core condition of data center development rather than a voluntary sustainability commitment. Policy options include requiring 24-hour carbon-free energy procurement, stricter permitting conditions for backup generators, and enhanced monitoring of generator testing and emergency operation.

On environmental justice: New Jersey's Environmental Justice Law requires cumulative impact assessments for certain covered facility categories in overburdened communities; however, data centers are not automatically included in these categories unless they trigger specific permit thresholds, such as qualification as a major air source. The screening results presented here suggest that expanding environmental justice review requirements to include large data center projects above a specified power capacity threshold could improve transparency and ensure meaningful community participation in siting decisions.

Limitations

Several limitations should be noted. First, the facility inventory is based on publicly disclosed information from DatacenterMap, Baxtel, and company disclosures; undisclosed or enterprise-only sites are not captured. The inventory is best understood as a conservative major-facility compilation for scenario analysis purposes rather than a complete statewide census.

Second, emissions estimates use EPA eGRID average output emission rates rather than short-run or long-run marginal emission rates. These are appropriate for average-grid attributional screening but would not be appropriate for decision-analytic assessments of the emissions consequence of incremental load growth, which would require marginal rate estimates and a different modeling framework.

Third, this study does not include facility-specific cooling-technology data, direct metered electricity consumption, or tract-level matched environmental justice comparisons. Water use, while an important environmental dimension of data center operations, could not be estimated or analyzed from the available data and is acknowledged here as explicitly outside the present empirical boundary. Future work should incorporate facility-specific cooling configurations and water consumption data where these become available.

Fourth, the environmental justice analysis is screening-level and does not constitute a statistically tested assessment of siting disparities. A more rigorous analysis would geocode every facility boundary, extract host-tract and adjacent-tract EJ indicators, and test differences against matched controls using nonparametric or regression methods.

Fifth, cross-case comparisons with Northern Virginia and Central Ohio are based on published studies and regional statistics rather than facility-level data constructed with identical methods, and should be interpreted as contextual benchmarks rather than direct equivalents.

CONCLUSION

This study demonstrates that the expansion of AI-driven data centers in New Jersey represents a potentially significant shift in the state's energy and environmental landscape. By translating publicly disclosed infrastructure capacity into system-level electricity demand, emissions, and health impact estimates under transparent screening assumptions, the findings indicate that data centers could become a major component of statewide energy consumption under full disclosed build-out — potentially accounting for approximately 11.6–15.5% of New Jersey's annual retail electricity sales, compared with approximately 2.8–3.7% under currently installed capacity.

The screening analysis does not establish exact future outcomes and should not be interpreted as a causal or predictive model. What it does demonstrate is that the scale of disclosed data center development in New Jersey is large enough to warrant advance policy planning on tariffs, transmission, cleaner incremental electricity supply, backup-power permitting, and tract-level environmental justice review.

Importantly, the modeled health damages and emissions burdens are not projected to be evenly distributed: the concentration of planned hyperscale capacity in Cumberland, Union, and Gloucester counties corresponds to the largest estimated impacts in the COBRA screening, raising equity concerns that deserve closer scrutiny.

Comparative analysis with Northern Virginia and Central Ohio demonstrates that the outcomes of rapid data center growth are shaped significantly by policy decisions. Ohio's development of data-center-specific tariff frameworks and Virginia's experience with grid reliability events both offer instructive precedents. Proactive regulatory attention to cost allocation, emissions disclosure, environmental justice review, and incremental clean energy supply can help ensure that AI infrastructure expansion does not come at the expense of public health, grid reliability, or social equity.

The most useful immediate step is to improve disclosure: facility-level capacity, cooling configuration, generator permit details, hourly load expectations, and host-tract environmental indicators should all be reported more transparently before the largest proposed projects proceed. More rigorous tract-level environmental justice analysis, incorporating geocoded facility boundaries and matched census-tract comparisons, is recommended as the priority next step for future research.

ACKNOWLEDGEMENTS

The author would like to sincerely thank Professor Shaolei Ren of the University of California, Riverside for his guidance and support as an advisor throughout this research. His insights on environmental impact modeling, particularly the application of EPA COBRA and energy system analysis frameworks, were instrumental in shaping the methodological approach of this study. The author also acknowledges the availability of publicly accessible datasets, including EPA eGRID, COBRA, and EJScreen, which made this research possible.

CONFLICT OF INTEREST

The author declares that there are no conflicts of interest related to this work.

REFERENCES

1. International Energy Agency. Electricity 2024: Analysis and Forecast to 2026. Paris: IEA; 2025. Available from: https://www.iea.org/reports/electricity-2024 (accessed 2026-04-15).
2. Shehabi A, Smith SJ, Masanet E, Koomey J. United States Data Center Energy Usage Report. Lawrence Berkeley National Laboratory; Report No. LBNL-2001441. 2024. Available from: https://eta.lbl.gov/publications/united-states-data-center-energy (accessed 2026-04-15).
3. Uptime Institute. Global Data Center Survey Results 2023. New York: Uptime Institute; 2023. Available from: https://uptimeinstitute.com/resources/research-and-reports/uptime-institute-global-data-center-survey-resu lts-2023 (accessed 2026-04-15).
4. Capgemini Research Institute. Generative AI and the Evolving Role of Enterprise. Paris: Capgemini; 2023. Available from: https://www.capgemini.com/insights/research-library/generative-ai/ (accessed 2026-04-15).
5. Stokel-Walker C. AI chatbots are the environmental villain you didn't know you had. The Guardian. 2023 Aug 1. Available from: https://www.theguardian.com/technology/2023/aug/01/techscape-environment-cost-ai-artificial-intelligen ce (accessed 2026-04-15).
6. O'Brien I. Data center emissions probably 662% higher than big tech claims. Can it keep up the ruse? The Guardian. 2024 Sep 15. Available from: https://www.theguardian.com/technology/2024/sep/15/data-center-gas-emissions-tech-companies-underre porting (accessed 2026-04-15).
7. Han Y, Wu Z, Li P, Wierman A, Ren S. The Unpaid Toll: Quantifying and Addressing the Public Health Impact of Data Centers. arXiv preprint arXiv: 2412.06288. 2024 Dec 9. Available from: https://arxiv.org/abs/2412.06288 (accessed 2026-04-15).
8. Dodge J, Prewitt T, Tachet des Combes R, Odmark E, Schwartz R, Strubell E, et al. Measuring the carbon intensity of AI in cloud instances. In: Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22). New York: ACM; 2022; p.1877–1894. doi: 10.1145/3531146.3533234.
9. Crawford K. Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence. New Haven: Yale University Press; 2021. ISBN: 978-0300209570. https://doi.org/10.12987/9780300252392
10. Mytton D. Data centre water consumption. *npj Clean Water*. 2021; 4 (1): 1. doi: 10.1038/s41545-021-00101-w.
11. Consumer Federation of America. Data Centers and Consumer Impacts: Rate Design and Cost Allocation Concerns. Washington DC: CFA; 2025. Available from: https://consumerfed.org (accessed 2026-04-15).
12. Data Center Dynamics. Renewable Energy Claims and Data Center Transparency. Available from: https://www.datacenterdynamics.com (accessed 2026-04-15).
13. Roberts D. The Problem with Renewable Energy Credits. Volts [newsletter]. Available from: https://www.volts.wtf (accessed 2026-04-15).
14. Center for Biological Diversity. Data Centers and Environmental Impact. Tucson: Center for Biological Diversity; 2025. Available from: https://www.biolo gicaldiversity.org (accessed 2026-04-15).
15. Virginia Joint Legislative Audit and Review Commission (JLARC). Data Centers in Virginia. Report to the Governor and General Assembly of Virginia. Richmond, VA: JLARC; 2024. Available from: https://jlarc.virginia.gov/landing-2024-data-cen ters.shtml (accessed 2026-04-15).
16. PSE Healthy Energy. Pollution, Health, and Environmental Justice Impacts of Backup Diesel Generators at Data Centers in Northern Virginia. Oakland: PSE Healthy Energy; 2023. Available from: https://www.psehealthyenergy.org (accessed 2026-04-15).
17. Schneider Electric. Data Center Market Overview: Q4 2024. Rueil-Malmaison: Schneider Electric; 2024. Available from: https://www.se.com (accessed 2026-04-15).
18. Columbus Dispatch. AEP Ohio proposes new rate structure for data centers. Columbus Dispatch. 2025. Available from: https://www.dispatch.com (accessed 2026-04-15).
19. Public Utilities Commission of Ohio (PUCO). In the Matter of the Application of Ohio Power Company.

Case No. 23-1651-EL-ATA. Columbus, OH: PUCO; 2025. Available from: https://dis.puc.state.oh.us (accessed 2026-04-15).

20. New Jersey Senate. Senate Bill S731: An Act Concerning Electric Utilities and Large Load Customers. 221st Legislature. Available from: https://www.njleg.state.nj.us (accessed 2026-04-15).
21. U.S. Environmental Protection Agency. Emissions & Generation Resource Integrated Database (eGRID) 2022. Washington DC: EPA; 2024. Available from: https://www.epa.gov/egrid (accessed 2026-04-15).
22. U.S. Environmental Protection Agency. Co-Benefits Risk Assessment (COBRA) Health Impacts Screening and Mapping Tool, Version 5.0. Washington DC: EPA; 2023. Available from: https://www.epa.gov/cobra (accessed 2026-04-15).
23. U.S. Environmental Protection Agency. EJScreen: Environmental Justice Screening and Mapping Tool, Version 2.3. Washington DC: EPA; 2024. Available from: https://www.epa.gov/ejscreen (accessed 2026-04-15).
24. New Jersey Department of Environmental Protection. Environmental Justice Mapping, Assessment and Protection (EJMAP) Tool. Trenton, NJ: NJDEP; 2024. Available from: https://dep.nj.gov/ej (accessed 2026-04-15).
25. U.S. Energy Information Administration. New Jersey Electricity Profile 2023. Washington DC: EIA; 2024. Available from: https://www.eia.gov/electricity/state/newjersey/ (accessed 2026-04-15).