

[Rania Masri's June 9th presentation notes in Greensboro, NC.](#) These notes support slide 2. Slides are [here](#).

[NCEJN](#)

[Note: these pages include expanded information and source material on what was presented during the presentation.] Please also see the tab (in the left under 'document tabs' to go directly to the section that interests you.) Please reach out to Rania if you have any questions - or to join the NC Data Center Network. Email: rania@ncejn.org

=====

What is an AI-Data Center?

In its simplest form, a data center is a physical facility that houses and runs large computer systems. Traditional data centers optimize for general workloads with CPU-based servers and balanced storage, while AI data centers specialize in GPU/TPU hardware to handle the massive parallel processing requirements of machine learning training and inference.

A CPU (Central Processing Unit) is the main processor that handles general computing tasks and runs the computer's operating system. A GPU (Graphics Processing Unit) is specialized hardware designed for parallel processing, making it much faster for tasks like AI training and graphics rendering that can be broken down into many smaller calculations at once.

In other words, traditional data centers store information and assist with general computing tasks. Traditional data centers power everyday services like your email and online banking - when you check Gmail or log into your bank account, you're using traditional data centers that process straightforward transactions and store your data in structured databases.

AI data centers, using Large Language Models, are used to train AI and create content - using facial recognition, for example. Or ChatGPT.

Sounds benign. But it isn't.

"The creation of contemporary AI systems depends on exploiting energy and mineral resources from the planet, cheap labor, and data at scale. It is an extractive industry"

Let's explore how.

Energy Consumption

Data centers accounted for 4% of total U.S. electricity use in 2024. Their energy demand is expected to more than double by 2030.

In other words, based on the current pace of AI computational infrastructure expansion, we would need to put as much energy on the global grid as what is consumed by two to six times the energy consumed annually by the state of California, and that will mostly be serviced by fossil fuels. We're already seeing reporting of coal plants with their lives being extended. They were supposed to retire, but now they cannot, to support this data center development. We are seeing methane gas turbines, unlicensed ones, being popped up to service these data centers, as well.

The number of US data centers more than doubled between 2018 and 2021 and, then doubled again by 2025. Early in the AI boom, in 2023, US data centers consumed 176 terawatt-hours of electricity, roughly as much as the entire nation of Ireland, and that's expected to double or even triple as soon as 2028.

By 2030, data centers are expected to use 12 percent of US electricity consumption, triple the current 4 percent. Even in lower estimates, this exceeds the amount of electricity needed to power 26 million households.

This energy grab is already destabilizing our energy grid, with the North American Electric Reliability Corporation (NERC) citing data center growth as one of the greatest reliability challenges for US power grids increasing the risk of power outages and bringing higher risks of energy shortfalls to virtually the whole country. In July 2024, data center operators almost caused blackouts across Virginia as they prioritized their own facilities' continued operations following an equipment failure.

IN NC: 80% of Duke Energy's projected demand is for data centers

- from 2000 to 2025: electricity demand increased 7% in NC
- from 2025 to 2040 - projected electricity demand is between 16% to 60%; 80% of that demand is from these hyper scale data centers

Duke Energy has received tax breaks for three data centers that it owns. And the utility is proposing an 18 percent rate hike on its North Carolina customers and reported \$4.9 billion in gross profits last year.

The billionaire financier behind the Deep River Data Project in Lee County wrote in 2024 - "The anemic demand for natural gas is linked to a near-absence of new electric power demand. There is a nearly universal consensus that the demand for fossil fuels will decrease over the next decade." Because of data centers, he writes, "electricity demand is now rising on a heretofore unknown scale" Currently, Lee County has nine solar farms that generate more than enough power for every household in the county. Without AI data centers, there is no reason to build out methane gas infrastructure in Lee County.

Let's zoom back out → According to a UN University report, released this month, - global data centers used 448 trillion watt-hours of electricity last year – more than all but 10 countries in the world. That electricity produced 208 million tons of carbon dioxide (the same as Argentina) and producing that much energy consumed about 1.20 trillion gallons of water. If data centers were a country, the country would be projected to rank 6th highest in power use in 2030, and would produce nearly 440 million tons of CO2

What about using renewable energy? That's what some are advocating.

- Right now, fossil fuels power the vast majority of data centers, contributing to their pollution, climate harms, and water use (generating electricity with fossil fuels consumes exponentially more water than doing so with renewable energy). Some developers are promising to power data centers with renewables. But that doesn't solve the problem.
- We need renewable energy everywhere. However, powering the massive energy demand of data centers with renewables, without limiting fossil fuel development elsewhere, will push dirty energy into other industries. We need to deploy renewable energy to meet all of our electricity needs, not just the ambitions of Big Tech.
- As data centers suck up the available supplies of renewable energy infrastructure, the cost of renewable energy will also go up for other sectors of the economy. Trump administration policies that threaten progress for renewable energy will only worsen this problem.
- Tech giants will be all too happy to greenwash themselves with renewable energy, while pretending the broader impacts of data centers don't exist.
- Data centers have to run 24/7, so they cannot actually run on just renewables. Solar does not actually work when we do not have sufficient enough energy storage solutions for that 24/7 operation.
- Some are advocating for nuclear energy instead. What? Nuclear!

Imagine, though, if all that energy were actually devoted to our transition away from fossil fuel usage and into solar and wind, a transition grounded in decreasing our consumption!

But wait! Even if the materials for renewable energy (wind and solar, and not nuclear) were available, we cannot have a livable planet with these AI data centers! Do we have enough water?

Water

Water is most intensively used for training AI → According to the UN University report: The water footprint associated with training GPT-4 was about 600 million liters, enough to meet the minimum annual domestic water needs of 81,000 people in Sub-Saharan Africa

AI operations likely consume between 312.5 billion and 764.6 billion liters of water per year—bracketing and potentially surpassing the 446 billion liters of bottled water people drink globally.

How is water used:

Data centers also use massive amounts of water for cooling their servers, which run super-hot. By 2028, U.S. data center water needs (just for cooling) could be as high as the indoor needs of 18.5 million U.S. households. Direct cooling isn't the only way data centers drain water supplies.

Nearly three-quarters of a data center's water use comes from generating electricity. And that footprint is several hundreds of times higher for cooling processes powered by fossil fuels and nuclear compared to renewables.

A single data center can consume as much water as a city of 50,000 people

Nevertheless, despite these water needs, roughly two-thirds of data centers built since 2022 are in water-stressed regions, raising concerns of worsening water scarcity problems. One Georgia county reported that taps ran dry after Meta began building a data center. Data centers also have implications for our water bills: in that same Georgia county, water rates will rise by 33% in two years, compared to the typical 2% annual increase.

In NC: For the past five years, NC has been experiencing alarming drought conditions. Currently, 97 of the state's 100 counties are in drought!

Land Loss

Massive data centers are reshaping how land is used in many communities. In 2024, the average data center site covered about 224 acres or 0.35 square miles — roughly the size of 450 football fields. This represents a 144% increase since 2022.

Developers are increasingly seeking large parcels for multi-building campuses that allow for phased construction and future expansion. Hyperscale facilities require even more land, with the largest campuses exceeding 1,000 acres, or about 1.6 square miles, according to a recent report from the Wisconsin Farm Bureau.

This often intersects with agricultural and rural land uses. In Wisconsin, where about 50 data centers are located, residents have concerns that rezoning for data center development could accelerate the loss of productive farmland, which is already declining due to development pressures, land consolidation and competition. Similar tensions have surfaced across rural communities in Ohio, Indiana and Virginia. In Ohio, a proposal to rezone 300 acres of farmland to support data center development sparked such widespread opposition that the landowners ultimately withdrew their application.

Public Money

14 - from 36 - states and scores of localities across the U.S. fail to disclose how much revenue they lose to data center tax abatement programs. Yet such losses are known to be soaring in states that do disclose, with three states, Georgia, Virginia, and Texas, already losing \$1 billion or more per year.

In NC: There are 3 state sales and use tax exemption programs for data centers

1. Electricity and equipment exemption for data centers
2. Electricity and business property for internet data centers
3. Computer software for data centers

There is a lack of transparency surrounding the cost of these programs, but the Governor's Energy Task Force estimates that from Program #1 alone:

- NC is currently losing about \$50 million annually
- If all projects that have been announced are completed, NC will lose:
 - between \$1.5 billion and \$2.3 billion during the construction phase
 - about \$400 million annually – a 700% increase

Local governments have also granted property tax exemptions to data centers

- No comprehensive database of these tax exemptions exists
- There is a lack of transparency about how much local governments are losing

The 10 year deal Apple made with the Town of Maiden and **Catawba County**

- 50% property tax exemption on land and buildings
- 85% property tax exemption on servers and equipment
- Estimated cost over 10 years: \$20 million
 - This is just a single data center – Catawba has reached the same deal with Microsoft

Google's data center in **Lenoir** paid only \$5 million in county taxes over 10 years but received \$73 million back in tax rebates—14 times what the county received. The city also abated \$51 million in property taxes over eight years (2017–2024).

AND WHAT DO DATA CENTERS GIVE BACK TO THE COMMUNITY?

Air and water pollution

By 2030, air pollution from data centers alone could cause an additional 600,000 asthma cases annually in the US and an estimated 1,300 premature deaths.

As the electricity grid struggles to keep up with the immense demand of AI data centers, companies are increasingly turning to dirty behind-the-meter power sources, further exacerbating public health impacts. The diesel generators frequently used for backup power release additional pollutants and toxins that cause asthma, cancer, heart attacks, and cognitive decline. Nitrogen oxide from gas turbines is making it hard for nearby residents to breathe: In the Boxtown neighborhood of Memphis, for example, xAI's supercomputer facility operates 35 turbines without pollution controls or adequate permits.

Water pollution - There are a lot of concerns over specific contaminants that are introduced to the water during the cooling process,

And PFAS → Such servers require cooling systems, semiconductors, and fire suppressants—all sources of PFAS forever chemicals. Chemours is one of the **largest global producers of PFAS** and one of the biggest suppliers of PFAS materials for **data center coolants** and **semiconductor production** (semiconductors are used to make the microchips deployed in data centers in vast quantities). The company aims to expand production of its PFAS materials, in particular resin for **semiconductor manufacturing**, to **accommodate the growth of data centers** across the country. This will likely have repercussions on the environment. In North Carolina, airborne PFAS emissions from Chemours operations have contaminated more than **7,000 drinking water wells**,

according to researchers.

Data centers increase PFAS pollution directly and indirectly. The chemicals are needed in the centers' operations – such as its cooling equipment – which almost certainly leads to some on-site pollution. The use of PFAS by data centers is under-reported and under-studied.

Noise:

What is the cause?

- “Data centers generate significant noise pollution primarily from diesel generators and Heating, Ventilation, and Air Conditioning (HVAC) systems, with internal noise levels reaching up to 96 A-weighted decibels (dBA)—well above the 85 dBA threshold considered harmful to hearing [Source is [here](#)].

What are the consequences?

- “Continuous operation of cooling systems and generators produces persistent noise exposure near data centers, creating a direct environmental stressor. Short-term effects include sleep disturbance and annoyance, while long-term exposure is associated with cardiovascular, cognitive, and mental health impacts (Babisch et al., 2013; Thompson et al., 2022).
- The health impacts of noise pollution include long-term risks such as heart disease, permanent hearing loss, and tinnitus; chronic effects such as hypertension and endocrine disruption; and acute effects like decreased sleep quality and temporary hearing changes (Babisch et al., 2013; Hammer et al., 2014).
- Noise pollution has also been linked with an increase in mental health concerns, such as stress-related disorders and increased annoyance, particularly in sensitive individuals (Meerlo et al., 2008; Sandrock et al., 2008).”

Who is impacted?

- Noise pollution from data center construction, generators, heating and cooling systems, and energy sources can also result in negative health consequences for workers and nearby residents (including headaches, stress, and disrupted sleep, which in turn lead to cognitive impairment and cardiovascular risks), as well as impacts on local wildlife. (source: <https://doi.org/10.1016/j.eehl.2025.100157>; Kelly Richardson, “Understanding the Impact of Data Center Noise Pollution,” Tech Target, December 3, 2024.)
- Impact of noise on property values: “adjacent residential property values can be reduced by 5 to 20 percent when industrial facilities produce sustained audible noise above local ordinance limits” <https://www.lsars.com/data-center-noise>

- One form of sound pollution is called **infrasound**, which is inaudible to humans but can be felt, and some claim it causes headaches, insomnia, nausea, and anxiety. [Heatmap Plus](#) reports that this is the phenomenon of frequencies so low they're inaudible to humans. Nevertheless, some people can feel it, and there have been claims linking them to various negative health effects such as headaches, insomnia, nausea, and anxiety. Infrasound and its effects need further study, but it's one of the issues local governments have been raising as they place a moratorium on data center projects. See this video: https://youtu.be/zy_ctHNLan8 ([Source](#))
- "Communities report **low-frequency hum** from AI data centers causing headaches and insomnia, prompting moratoriums and calls for more research. [Source](#)

The long-term impacts of both noise and light pollution include hearing loss, stress, insomnia, and decreased quality of life.

- Noise pollution is regulated at the local and state levels through zoning ordinances, but for a time in the 1970s, the Environmental Protection Agency oversaw noise pollution and conducted noise-control investigations. It still has the legal authority to do so.
- Because there is a lack of reliable data from sound level meters and because most county or community noise ordinances are written to address noisy block parties (rather than data centers), most noise complaints go nowhere, leading to frustration and lower property values.

Heat island effect

They are creating "[heat islands](#)," warming the land around them by up to 16 degrees Fahrenheit, and making life hotter for more than 340 million people.

They found surface temperatures increased by an average of 3.6 degrees Fahrenheit after a data center started operations. In extreme cases, nearby temperatures increase by up to 16.4 degrees Fahrenheit. [Source](#)

Runoff and flooding

headlines from May 25 → Data center construction site runoff floods Mason County neighborhood // Data Center Construction Linked to Flooding Damage in Nearby Neighborhood in West Virginia

Increased housing costs

[Housing costs can skyrocket. While the city or county's ability to allow more homes to be built decreases.](#)

And more: The physical lifecycle of AI hardware presents a growing crisis. AI infrastructure could generate up to 2.5 million metric tons of e-waste annually by 2030, equivalent to discarding nearly 250 Eiffel Towers every year. The environmental impacts of e-waste are primarily linked to its toxic components, such as lead and mercury, which have the potential to contaminate water and soil and pose serious environmental risks to humans and non-human

And all this for what purpose: [Harms of “AI”]

Loss of Jobs

Experts say [AI will disproportionately impact Black and brown workers through displacement.](#)

400 million workers could be displaced because of AI As AI evolves, it could displace 400 million workers worldwide. A McKinsey report predicts that between 2016 and 2030, AI-related advancements may affect around 15% of the global workforce.

40% of employers expect to reduce their workforce where AI can automate tasks.

Pricing and wage algorithms

Algorithmic wage discrimination develops a labor market where workers perform the same services, with the same skill, for the same company, but receive different hourly wages Workers’ pay is increasingly shaped by opaque algorithms and artificial intelligence systems, shifting compensation decisions away from human managers, clear legal standards, and collective bargaining. This phenomenon—known as algorithmic wage discrimination or surveillance pay²—was first documented in app-controlled ride-hail and food-delivery work.

People who work longer hours are paid less per hour.

According to 2023 [research](#) by Veena Dubal, companies collect data on each worker, including where they live, what hours they normally work, and what their typical earnings target is, and [“use that data to personalize and differentiate wages for workers in ways unknown to them.](#)

Now, it is spreading to a range of other industries and services. Traditional employers in industries including health care, customer service, logistics, and retail are now using automated surveillance and decision-making systems to set compensation structures and to calculate individual wages [Source](#):

In the third quarter of 2025, the "labor share" of Gross Domestic Product in the U.S. [decreased to 53.8%](#), according to data from the Bureau of Labor Statistics. This was the lowest level seen since the Bureau began recording labor share in 1947. One likely factor: the rise of algorithmic pay. In recent years, major companies have begun to use artificial intelligence (AI) algorithms to set wages for some workers.

Denial of rights:

[92 million low-income people in the U.S. states—everyone whose income is less than 200 percent of the federal poverty line—have some basic aspect of their lives decided by AI.](#) Government officials, business leaders, and other people in positions of power regularly use AI and related technologies to wrongly [cut Medicaid](#) and [other public benefits](#), deny people [housing](#), [employment](#), or [credit](#), take [kids from loving parents](#), intensify [domestic violence](#) and [sexual abuse or harassment](#), label and mistreat middle- and high-school kids as likely [dropouts](#) or [criminals](#), and [falsely accuse](#) Black and brown people of crimes, all at scales and speeds impossible with analog methods.

Fake porn, scams, and false information

[Tristan Harris](#), co-founder and executive director of the Center for Humane Technology, warns us about Generative AI. He says: "[Floods of deepfakes and AI-generated content overwhelm our information environment.](#)" Remember: 90% of online content will be generated by AI this year. Conservatively, 62% of internet traffic is non-human.

Even the Department of Homeland Security is speaking out on this. ["Deepfakes, an emergent type of threat falling under the greater and more pervasive umbrella of synthetic media, utilize a form of artificial intelligence/machine learning \(AI/ML\) to create believable, realistic videos, pictures, audio, and text of events which never happened"](#)

Increased surveillance

[Palantir](#) is the software powering intelligence agencies, health insurance companies, private hospitals and defense department systems, consolidating your data to build algorithmic predictive models. The company wants to monopolize and weaponize surveillance to maximize profits for clients and increase the state's surveillance

capabilities—all at the expense of your privacy.

The American Civil Liberties Union revealed that Palantir Technologies plays a central, enabling role in the mass deportation campaign carried out by ICE. Through four main systems built or operated by Palantir that underpin this effort: (1) ELITE system – identifies and maps targets for deportation using large datasets. (2) Investigative Case Management (ICM) – a massive data platform aggregating personal information across agencies to build cases. (3) ImmigrationOS – enhances tracking, prioritization, and logistics of deportations, including “self-deportation” monitoring. (4) AI-enhanced tip processing – uses AI to analyze and translate public tips. These tools collectively form the “backbone” of the deportation system, raising serious concerns about privacy, due process, and human rights.

USDA just gave Palantir a \$300 million no-bid contract to consolidate every farmer's records into one digital file. This was one day after a watchdog sued five agencies over Palantir surveillance systems ([American Oversight Sues for Records on Trump Admin Use of Palantir Tools Amid Growing Privacy, Data Collection Concerns Federal agencies should disclose whether Palantir tools are being used for surveillance.](#))

Genocide

See this infographic [Automating Genocide: Israel's Use of AI-Driven Warfare](#)

And more information [here](#): Reports on the ‘[Lavender](#),’ ‘[Gospel](#),’ and ‘[Where's Daddy](#)’ AI-enabled data processing systems developed and in use by the Israeli Occupation Forces (IOF) in their genocidal campaign against Gaza have caught widespread attention, prompting [journalists](#) to call Gaza the site of the first AI-powered genocide. [AI technology](#) was reportedly [first used](#) in Gaza during Israel's 11-day assault in 2021. During the ongoing genocide, for the first time, it is being used to kill Palestinians at an unprecedented level and at much faster rates. These three known [systems](#) identify “targets” for airstrikes based on Israeli mass surveillance records of Palestinians in Gaza that have been collected for years by the IOF under the racist framework of monitoring what they deem as “threats” to the Israeli regime. The ‘Gospel’ “recommends” buildings and structures to strike, while the ‘[Lavender](#)’ and ‘Where's Daddy’ systems “recommend” people to kill and track their location to determine when a strike should be administered.

AI was also used in the US/Israeli war on Iran.

Unknown harms

“Frauds and scams skyrocket,” Harris says. “AI increases hacking capabilities that makes critical infrastructure more vulnerable. AI enables bad actors to do dangerous new things with biology.”

Listening to Trisan Harris, I got scared. He is scared. He [says](#), “I used to be very skeptical when friends of mine in the AI community worried about sci-fi scenarios of AIs scheming, lying, or trying to escape. But unfortunately over the last few months, we’ve seen clear evidence of what was previously the realm of science fiction happening in real life. We’re now seeing frontier AI models [lie and scheme to preserve themselves](#) when they are told they will be shut down or replaced; we’re seeing AIs [cheating when they think they will lose](#) a game; and we’re seeing AI models unexpectedly attempting to [modify their own code to extend their runtime](#) in order to continue pursuing a goal. So we don’t just have a country of Nobel Prize geniuses in a data center—we have *a million deceptive, power-seeking, and unstable geniuses*. ... We’re currently releasing the most powerful, inscrutable, uncontrollable, omni-use technology we’ve ever invented, one that’s already demonstrating the self-preservation, deception and escape behaviors we previously thought only existed in science fiction movies, and we’re deploying it *faster* than we’ve deployed any other technology in history, under the maximum incentives to cut corners on safety.”

In an [interview](#), Harris spoke about “how AI could trigger a global collapse by 2027 if left unchecked, and how AI will take 99% of jobs globally and collapse key industries by 2030.”

Quote from Kate Crawford

“Artificial intelligence is not an objective, universal or neutral computational technique that makes determinations without human direction. Its systems are embedded in social, political, cultural, and economic worlds, shaped by humans, institutions, and imperatives that determine what they do and how they do it. They are designed to discriminate, to amplify hierarchies, and to encode narrow classifications. When applied in social contexts, such as policing, the court system, health care, and education, they can reproduce, optimize and amplify existing structural inequalities. This is no accident: AI systems are built to see and intervene in the world in ways that primarily benefit the states, institutions, and corporations that they serve. In this sense, AI systems are expressions of power that emerge from wider economic and political forces, created to increase profits and centralize control for those who wield them.” - [Kate Crawford, Atlas of AI](#) // Also see the [Anatomy of AI](#)

ADDITIONAL INFORMATION / Answers to Frequently Asked Questions

What are the most dominant greenwashing statements?

Please see this article by Food and Water Watch on the [7 Greenwashing Lies](#).

Is the influx of data center development an opportunity?

The city should not assume data center development is an “opportunity” until developers prove the public benefits outweigh the public costs. Data center development should be evaluated as high-impact industrial infrastructure, not presumed to be a public opportunity. The burden should be on developers to demonstrate a measurable community benefit, full cost coverage, and no added environmental or utility burdens on residents.

The only recorded ‘benefit’ of data center development to communities is a negotiated property tax for the county or the city. Yet, when weighed against the state sales and usage tax incentive, then the city or county would be getting less than pennies to the dollar. When further weighed against the massive cost of energy expansion, water consumption, air and water pollution, land transformation, and increased surveillance from the AI itself, then even these meager pennies dissipate.

Would an information-sharing forum help to level the playing field?

Unfortunately, an impact forum is not enough to level the playing field, since it does not in any way guarantee transparency or public disclosure. Access to developer-provided data does not equal power, expertise, or accountability. Developers can selectively disclose information, hide risks, or claim confidentiality. Council members and residents need more than access to data.

Instead, we recommend banning NDAs, mandating public disclosure at least 90 days before any public meeting, and requiring independent studies (which can be funded by developers but not conducted by them) as part of any process before approval. The city, and its residents, need independent analysis, plain language documentation, public disclosure requirements, and the ability to challenge developer assumptions. More here: <https://datacenters.ainowinstitute.org/issue-area/transparency/>

Developers have more money, data, consultants, lawyers, and access to decision makers. Residents enter with lived experience and fewer resources. We should not pretend that all parties have equal power or equal exposure to risk. The goal should be

accountable decision making that gives residents, public health experts, and environmental justice communities meaningful power to shape the rules.

Developers already have technical consultants, lawyers, engineers, and utility relationships. Starting from developer based information is starting from a place of pro-data center bias. Residents and council members need independent analysis, disclosure requirements, and the authority to say no.

“Dialogue” without enforceable policy can slow opposition while projects continue moving forward leaving communities vulnerable to impacts that are locked in long before the dialogue is complete.

What about incentivizing developers to invest in local renewable energy and battery storage that supports grid resilience?

The first problem is the word “incentivized.” If the city is serious about protecting residents, then local clean energy and storage should not be optional or treated as a bonus. It should be a binding condition of approval. Developers should not receive public benefits, tax incentives, infrastructure support, or expedited permitting unless they can prove that their project will not increase fossil fuel generation, shift costs to ratepayers, or otherwise consume clean energy that would otherwise go to homes, schools, and public infrastructure.

24/7 carbon-free matching is not the same as local impact reduction. This claim may sound good on paper, while still leaving communities exposed to higher rates, new transmission burdens, expanded gas pipelines, and diesel backup emissions. A developer can match electricity use with clean energy contracts or offsets, while the actual facility pulls from the local grid that relies on fossil fuels during peak generation. If the matching happens elsewhere, or through contracts that do not directly relieve grid, the community can still carry the infrastructure and pollution burdens while the company gets to claim good climate stewardship.

The city should require data center developers to demonstrate, before approval, that 100% of their projected energy demand will be met with new, local carbon free energy and storage that does not raise costs for existing ratepayers, does not delay the retirement of fossil fuel generation, does not trigger new gas infrastructure, and provides measurable resilience benefits to the surrounding communities.

This is all based on incentives, not requirements, which is weak. Any renewable energy requirements should also specifically be additive/additional, so they do not take away from other projects.

The issue is not whether a data center can buy a cleaner-looking power portfolio. The issue is whether residents are left paying for the grid strain, infrastructure buildout,

pollution risk, and opportunity cost of handing scarce clean energy capacity to one of the most electricity- intensive industries in the country.

What about closed-loop cooling systems or liquid cooling technologies that virtually eliminate the consumption of municipal water?

Liquid cooling and closed-loop systems can reduce direct municipal water consumption, but they do not “virtually eliminate” the broader water footprint of a data center.

They can shift the burden from direct water withdrawal to higher electricity demand, and electricity generation can carry its own water footprint depending on the power source. The water-energy tradeoff matters. Data centers often rely on a mix of air cooling, evaporative cooling, and liquid cooling, and the most water-efficient system is not always the lowest-energy system.

The phrase also lets developers hide behind technology categories rather than measurable outcomes. “Liquid cooling” is not a public protection standard. A meaningful requirement would specify maximum annual water withdrawal, peak-day demand, source water, wastewater discharge characteristics, drought restrictions, water-use effectiveness, and independent verification.

The city should not accept “closed-loop” or “liquid cooling” as a proxy for low impact. Developers should be required to disclose projected and actual water use, source water, blowdown or discharge volumes, chemical treatment practices, indirect water impacts from electricity demand, and drought-condition operating plans. The standard should be measurable performance, not a promise that a particular technology will solve the problem.

This is especially important because “virtually eliminate municipal water” can be technically true in one narrow sense while misleading in the broader public-interest sense. A facility could reduce municipal water demand but increase grid demand, trigger more generation, worsen peak load, or rely on water-intensive power elsewhere.

What about green buffers that absorb sound?

Green buffers are useful, but they are not a substitute for proper siting, industrial setbacks, noise limits, heat controls, and enforceable design standards. Trees and bioswales can reduce stormwater runoff and help mitigate heat islands; EPA recognizes

vegetation as a heat-island and stormwater strategy. But that does not mean trees can meaningfully offset the full operational footprint of a large data center, especially one with constant cooling equipment, large impervious surfaces, backup generators, truck traffic, and high energy demand.

The problem is scale. A “green buffer” can soften the edge of a facility, but it cannot neutralize the facility’s industrial impacts. It should not be used as visual mitigation for an otherwise inappropriate site. If a project requires a buffer to make it tolerable next to homes, schools, parks, or already overburdened neighborhoods, the city should ask whether that site is appropriate at all.

Green buffers should be treated as a minimum design requirement, not as mitigation sufficient to justify siting a data center near burdened communities. Tree canopy and bioswales cannot compensate for excessive energy demand, diesel generator emissions, industrial noise, heat discharge, stormwater runoff, or public infrastructure strain. The city should require buffers in addition to — not instead of — setbacks, noise caps, generator restrictions, heat mitigation standards, stormwater performance requirements, and cumulative impact review.

Green buffers stops far short of stronger protections, such as prohibiting data centers [in certain zones](#), or establishing more [robust setback requirements](#)

Also, “absorb sound” needs caution. Vegetation can provide some sound attenuation, but it is not a reliable primary noise-control strategy for industrial-scale fans, chillers, transformers, and backup generators. Noise mitigation should require acoustic modeling, enforceable decibel limits at property lines, equipment enclosure standards, restrictions on generator testing hours, complaint response protocols, and penalties.

What about Community Benefit Agreements (CBAs)?

Unfortunately, CBAs do not “ensure that data center developers contribute.” We know of numerous examples, including in NC, where CBAs are not enforced. CBAs do not guarantee any of what they include, but are individually negotiated and only as strong as their conditions. We recommend pursuing more proactive protections such as limitations in zoning (if bans are not an option) and if CBAs are asked for by community, they should be required within a CBA ordinance (Community Benefit Ordinance) or as part of the approval process -- but NOT be the only vehicle to hold developers accountable.

Community benefits should not be framed as offsets for avoidable harm. The first obligation is prevention. Any community benefit agreement should be legally binding, independently monitored, and additive to strict environmental, energy, water, and

land-use requirements. Benefits cannot replace compliance, and they should not be used to buy public acceptance of a project that increases cumulative burdens.

CBAs can be useful, but only when they are enforceable, community-shaped, and tied to measurable outcomes. Brookings has emphasized that data center CBAs should be legally binding and include quantifiable data and tracking; [Good Jobs First](#) similarly notes CBAs can address generator protocols, notification systems, and pollution concerns.

What about asking data center developers to offset their environmental footprint?

Data centers should not be allowed to “offset” harms that could have been avoided through better siting, smaller scale, cleaner power, non-diesel backup, water limits, and enforceable operating conditions. Offsets often become a permission structure: the developer creates a burden, then funds a partial community benefit that does not actually remove the harm.

Air monitoring and weatherization grants may be useful, but they are not equivalent to preventing pollution, heat, water stress, grid strain, or utility cost increases. Air monitors can document harm; they do not prevent it. Weatherization can lower household energy burden; it does not erase the ratepayer, infrastructure, and environmental impacts created by a high-load facility.

What about ‘stifling growth’?

This uses the industry's framing that regulation equals obstruction. The city does not need to apologize for protecting its residents and resources. Requiring responsible siting, disclosure, and enforceable protections is not “stifling growth.” It is governance. If a project cannot meet basic public health, utility, water, and environmental standards, then the problem is not the regulation, it is the project.

Are there any members of Congress that are pushing legislation in DC against data centers?

Yes. In December of last year, Vermont independent Sen. Bernie Sanders endorsed a [national moratorium on data center construction](#). In March, Senator Sanders (I-Vt.) and Rep. Alexandria Ocasio-Cortez (D-N.Y.) also announced the [Artificial Intelligence \(AI\) Data Center Moratorium Act](#), legislation that would enact a reasonable pause to the development of AI to ensure the safety of humanity.

[Stop the Data Center Coalition](#) is a national coalition of organizations also calling for national moratorium

Is there an online site that would be informative about how many data centers are existing and in process of being built?

There are many sites, but not one comprehensive national site.

Here are some good resources:

- NC Data Center Activity [database](#)
- NC Data Center Locations [database](#)
- Who's building data centers in your backyard? [interactive map](#)
- Data Center Map [interactive map](#)
- Oil, Gas, and Petrochemical – FracTracker Alliance [databases](#)
- Data Center Laws and Regulations – We Build Progress [database](#)
- Data Center Impact [dashboard](#)
- Tracking American AI Data Center Buildout – Community Response and Legislative Action [map](#)
- Community Organizations Focusing on Data Centers [map](#)
- Data Center Policy [Database](#)

Please see this [link](#) for additional resources and toolkits.