

# **AI** **Data Center** **Ecology Report**

**Our Water, Power  
and Local Actions**

## TABLE OF CONTENTS

|                                                               |           |
|---------------------------------------------------------------|-----------|
| <b>- Introduction</b>                                         | <b>5</b>  |
| <b>01 - Finite Power Against Infinite Hunger</b>              | <b>9</b>  |
| <b>02 - What's In A Data Center?</b>                          | <b>14</b> |
| <b>03 - Who's Driving These Data Centers?</b>                 | <b>19</b> |
| <b>04 - Where Are These Data Centers Located?</b>             | <b>24</b> |
| <b>05 - AI's Power Life Cycle</b>                             | <b>28</b> |
| <b>06 - Estimating Data Center Demands</b>                    | <b>31</b> |
| Efficiency Metrics Vs Raw Resource Usage                      | <b>33</b> |
| Which Estimate Is Correct?                                    | <b>40</b> |
| <b>07 - Big Tech's Sustainability Facade</b>                  | <b>42</b> |
| Data Centers May Not Be Measuring Water Usage                 | <b>43</b> |
| Sustainability Report Breakdown                               | <b>46</b> |
| Comparing Water Use                                           | <b>49</b> |
| Conclusion                                                    | <b>51</b> |
| <b>08 - Our Water Systems</b>                                 | <b>53</b> |
| How Aquifers Grow And Deplete                                 | <b>55</b> |
| <b>09 - Dangers To Water Systems</b>                          | <b>62</b> |
| Impacts Of Severe Ground Water Depletion                      | <b>63</b> |
| <b>10 - U.S. Water Rights</b>                                 | <b>71</b> |
| Texas And Groundwater Conservation Districts                  | <b>75</b> |
| <b>11 - Hyperscale Life Cycles: Community Wins And Losses</b> | <b>78</b> |
| Site Selection                                                | <b>82</b> |
| Going Industrial                                              | <b>83</b> |
| Is It Over?                                                   | <b>90</b> |
| Tax Breaks                                                    | <b>91</b> |
| Power And Energy Contracts                                    | <b>94</b> |

## TABLE OF CONTENTS

|                                                                         |            |
|-------------------------------------------------------------------------|------------|
| Environmental Permits .....                                             | 99         |
| Conclusion .....                                                        | 105        |
| <b>12 - How To Spot and Stop a Data Center .....</b>                    | <b>106</b> |
| <b>13 - Legal Climate .....</b>                                         | <b>121</b> |
| Emergency Grid Regulation .....                                         | 122        |
| States Giving Billions In Tax Subsidies .....                           | 122        |
| Data Center Moratorium Bills Are Going Dead .....                       | 127        |
| <b>14 - Gigascale Projects Transforming The Us .....</b>                | <b>129</b> |
| #6 — “Ran Too Fast”: Project Ranier In New Carlisle, Indiana .....      | 132        |
| #5 — “Not Fair”: Fairwater In Mount Pleasant, Wisconsin .....           | 136        |
| #4 — “Gate To Ruin”: Project Stargate In Abilene, Texas .....           | 140        |
| #3 — “Colossal Problems”: Colossus From Xai In Memphis, Tennessee ..... | 145        |
| #2 — “Hyper-Risky”: Meta’s Hyperion In Richland, Louisiana .....        | 150        |
| #1 — “Project Bull”: Fermi America’s Project Matador .....              | 156        |
| Conclusion .....                                                        | 161 / 162  |
| <b>15 - Dangers To Living Things .....</b>                              | <b>163</b> |
| Air Pollution .....                                                     | 164        |
| Health Impacts .....                                                    | 166        |
| Data Heat Island Effect .....                                           | 168        |
| Noise Pollution .....                                                   | 172        |
| <b>16 - Ai Baked In .....</b>                                           | <b>175</b> |
| The Ai New Normal .....                                                 | 177        |
| Big Tech Ai Overview .....                                              | 178        |
| <b>17 - Settings You Should (Try To) Disable .....</b>                  | <b>180</b> |
| Ios (Apple Intelligence) .....                                          | 181        |
| Android (Gemini) .....                                                  | 183        |
| Samsung .....                                                           | 184        |

## TABLE OF CONTENTS

|                                                         |            |
|---------------------------------------------------------|------------|
| Computers .....                                         | 185        |
| Browsers .....                                          | 189        |
| Social Media .....                                      | 191        |
| <b>18 - Calculating A Model's Energy Costs .....</b>    | <b>196</b> |
| <b>19 - Guidelines For Economical Ai Use .....</b>      | <b>201</b> |
| <b>20 - The Purpose Behind Data Centers .....</b>       | <b>205</b> |
| Speculative Bubble .....                                | 206        |
| Ai War Deterrence .....                                 | 211        |
| Robotics .....                                          | 218        |
| Surveillance And Predictive Policing .....              | 223        |
| <b>21 - Solutions .....</b>                             | <b>227</b> |
| <b>22 - Conclusion: A Spiritual Tipping Point .....</b> | <b>236</b> |
| <b>Acknowledgements .....</b>                           | <b>239</b> |
| <b>About The Author .....</b>                           | <b>240</b> |
| <b>Ai Statement .....</b>                               | <b>242</b> |
| <b>Resources .....</b>                                  | <b>243</b> |
| Reports & Papers .....                                  | 244        |
| Web Links .....                                         | 251        |
| Our Water Systems .....                                 | 256        |
| Water Rights .....                                      | 256        |
| Data Center Project Lifecycle .....                     | 257        |
| Why .....                                               | 258        |





## Introduction

The waves of technological progress have been relentless in the past decade. We routinely use search engines, mobile phones, e-commerce, content streaming, and smart devices — products of the digital age.

Now most of these products and ecosystems are being embedded with AI, with or without our consent, supposedly to make our lives easier.

# These systems aren't free or cheap. We are all paying the unspoken cost as AI data centers guzzle our natural resources.

Our focus in this report will be AI data centers, the next generation of data centers that power modern AI systems. They are facilities that span millions of square feet and house tens of thousands of AI servers, cost billions of dollars, and are becoming so vast that they now use as much power and water as an entire metropolitan city or even more than an entire state.

It's only after the waves of new technology crash and recede that we notice the damage left in their wake, the deep marks of erosion on society and the psyche of the individual.

Given that AI is developing at such an incredibly rapid pace, we cannot ignore its downsides. If left unchecked, its very real dangers will have catastrophic consequences for our local communities and earth as a whole.

## Indeed, we risk a future where:

- Water and utility bills double, triple, or even quadruple
- Wells on private property run dry from groundwater depletion
- Rolling brownouts and blackouts and power grid instability are the norm
- Entire ecosystems suffer from drought

In this comprehensive report, “AI Data Center Ecology,” we share with you an educational foundation to discuss hyperscale data centers; from knowing how to calculate their power, to their legal lifecycle, to how they may impact groundwater systems.

Additionally, this report accompanies our ongoing effort to collect information on thousands of data centers on **AI Data Center Map (<https://aidatacentermap.org>)**.

On this map you can zoom into your town, and see the impact of nearby data centers on power and water demand. You can also overlay data centers across important landmarks like aquifers and powerlines.

## Our intention is to inform three kinds of readers:

- Individuals who want to understand how AI data centers operate
- Anyone who works with AI systems and wants to gain new insights on how to do so with care
- Citizens who want to organize locally and who seek an informed starting point

**We appreciate your initiative, and we trust you will learn something new.**

01

# Finite Power Against Infinite Hunger







Big tech execs at the signing

We start this report by looking at government & industry's doublespeak in response to public awareness of data center power & water consumption.

In March 2026, initiated by President Trump, the heads of major tech companies gathered at the White House to sign the Ratepayer Protection Pledge.

Signatories of this voluntary pledge were Amazon, Google, Meta, Microsoft, Oracle, and xAI. The pledge promises the companies shall build or buy new power supply, pay for infrastructure upgrades, invest in local jobs, and even offer their own power in times of scarcity.

On the very same day as the signing, Trump confirmed his intention to restore 100% bonus depreciation using the One **Big Beautiful Bill Act**.

This bonus depreciation could be applied towards taxes, enabling AI data centers and other industries to write off the full costs of machinery and equipment retroactively since January 2025.

For major industries, this is a huge cash flow advantage. Its effectively an interest-free loan from the government. According to Senator Elizabeth Warren's office; tech companies deducted \$152 billion from their taxes in 2022, the last year with 100% bonus depreciation in effect.

Bonus depreciation from the Big Beautiful Bill is permanent until a future change. Every cent that major companies put into equipment and infrastructure can be deducted from taxable income.

The Joint Committee of Taxation estimates that corporations would save around \$178B in taxes for 2026.



While consumers suffer from a 267% jump in wholesale electricity prices (Bloomberg) in areas located near data centers — the companies causing the price hikes can write off billions of dollars in data center development.

There's no silver lining in the ratepayer protection pledge, it doesn't outline concrete numbers or specifics on how much power should be produced, or the number of local jobs that should be created. And of course, no outline for enforcement or accountability.

Ironically, the 100% bonus depreciation exacerbates the core issue; it encourages the quick buildout of data centers which can go from blueprint to operational within a year. Powerplants and power infrastructure however, can't keep up with the rapid pace and scale of construction.


A **2025 study** by Open Energy Outlook indicates that “demand” for data centers could increase nationwide electricity costs 8% by 2030, with the weight being passed down to the consumer.

Our estimates are far more bleak. Based on data from our map, if the U.S successfully brings all planned data centers into operation by 2030, the current AI power demand will quadruple from roughly 50 to 200GW, which is 40% of the current U.S power supply.

The reality is that even if data centers “buy their own power”, electricity prices will soar, untethered and encouraged by bonus depreciations. There’s no way to protect ratepayers from this type of demand.

## Server Demand Raises Power System Costs by 2030

Increase in demand-weighted marginal operational costs across U.S. regions



Central and Northern Virginia hit with price increases of more than 25%

02

# What's In A Data Center?



The core work in Artificial Intelligence systems takes place primarily in AI data centers. Whether it's predictive AI or generative AI, the operation of AI models involves massive amounts of data and computing power.

## Hyperscale Data Centers

These are the largest class of data centers, typically using 100MW or more of power, equivalent to 80,000 U.S. homes. The largest hyperscalers announced have been in the 10GW+ range, meaning they use more power than 8 million homes, double the total power of some states.

## Conventional Data Centers Versus AI Data Centers

The extreme workloads of AI demand huge data centers that are much different from conventional or cloud computing data centers. While conventional centers power services like email, file storage, applications, and video streaming, AI data centers are busy training and operating models on huge swaths of the internet.

AI datasets can grow up to hundreds of terabytes in size and contain all of the content on a platform, such as Reddit, GitHub, and Wikipedia.

To process huge amounts of data quickly, AI data centers use clusters of high-density GPUs (graphics processing units), whereas conventional data centers use CPUs (central processing units).




For example, in 2020, OpenAI's GPT-3 model was trained using 10,000 NVIDIA V100 GPUs in a Microsoft data center for some 15 days in a row. The training used an estimated 1.29 gigawatt hours (GWh) of energy, enough to power a million homes for an hour.

Three years later, its successor GPT-4 required 40 times more energy: 50 GWh. That's the equivalent of powering 5,000 homes for an entire year. It's also equal to 0.02% of California's annual electricity consumption.

Training AI is only the first stage in the process. We'll discuss later how 60–90% of total energy costs come from “inference”—that is, from actually using the AI.



# Smaller Size, More Power



## A single AI rack can draw up to 1 MW

Footprint similar to a large refrigerator

By 2028, next-generation AI server racks are expected to reach power densities of up to 1 MW each, roughly comparable to the continuous electricity use of 800–1,000 average U.S. homes.



[aidatacentermap.org](http://aidatacentermap.org)

Compared to conventional data centers, the GPUs in AI data centers take up a small amount of space. As a consequence, each server rack in an AI data center demands far more power and generates more heat than is the case in a conventional data center.

For example, the NVIDIA H100, a commercial AI GPU, performs 989 trillion operations per second.

Today's modern AI server rack draws anywhere from 80 to 100 kilowatts (kW) of power, or 20 times the average data center rack in 2010.

In 2028, next-generation AI servers are expected to consume 1 megawatt (MW) each—enough energy to power 1,000 US normal-sized homes, yet packed into a space no bigger than a refrigerator.

|                                                | Conventional Data Center | Hyperscale AI Data Center |
|------------------------------------------------|--------------------------|---------------------------|
| Components                                     | CPU (Multi-core)         | GPU / TPU                 |
| Single Server Power                            | 0.5–2.5 kW               | 5–12.5 kW                 |
| Rack Power Usage                               | 5 kW                     | 80–100 kW                 |
| Avg. Facility Power Usage                      | 10–25 MW                 | 100–1,000 MW              |
| Facility Equivalent to Homes Powered In 1 Year | 2,500                    | 100,000+                  |

How does something so small consume so much energy? Each GPU inside an AI server has billions of transistors. NVIDIA's Blackwell GPU, for example, has more than 200 billion transistors. Extremely high and concentrated heat is produced when these transistors charge and discharge.

At full load, the heat generated from these chips is within one magnitude of the surface of the sun. Such tremendous heat must be cooled through systems that use clean water.

Thus, in areas where water is scarce, data centers threaten the entire ecosystem.



03



—

# Who's Driving These Data Centers?



# In 2025, 6 companies pledged \$2.4T on data center projects.

## They include:




As you can see, there's an AI race afoot. Egos as massive as the hyperscalers they build, pushing their companies to build the largest hyperscale data centers in the shortest amount of time. We find among those inflated egos (with their inflated promises):

“

We continue to believe AI is a once-in-a-lifetime reinvention of everything we know. The demand is unlike anything we've seen before

**Andy Jassy**  
Amazon



“

xAI will have more AI compute than everyone else combined in <5 years.

**Elon Musk**  
X



“

Oracle will build more cloud-infrastructure data centers than all of our infrastructure competitors combined.

**Larry Ellison**  
Oracle





“

Meta is planning to build tens of gigawatts this decade, and hundreds of gigawatts or more over time.

aiddatadentarmap.org



**Mark  
Zuckerberg**  
Meta

“

Theoretically, at some points, you can see that a significant fraction of the power on Earth should be spent running AI compute.

aiddatadentarmap.org



**Sam Altman**  
OpenAI

“

GDP growth in any place will be directly correlated to the cost of energy in using AI.

aiddatadentarmap.org



**Satya Nadella**  
Microsoft

And we're only just getting started describing the big plans the data center industry has for itself.


McKinsey & Company suggests another **\$5 trillion** will be required in data center capital investment by 2030. This off-the-charts figure may explain why the flagship AI sites are rapidly building new facilities year after year.



04

# Where Are These Data Centers Located?






The International Energy Agency projects that the United States will experience the fastest expansion of data centers globally, with China and Europe close behind.




Unlike other industries, data center sites tend to be located close to where people live, typically within a few miles of urban areas. Data center sites are also built in clusters to take advantage of shared power and water infrastructure. Their location puts huge demands on the water and energy grids that house them. The largest data centers under construction use the same amount of power as millions of homes.

**Data centre electricity consumption in household equivalents**



**Spatial concentration and urban proximity of various facilities**




IEA. CC BY 4.0.

Data centers tend to be geographically concentrated and located around cities. A 100 MW data centre can consume as much electricity as 100,000 households.



Data centers now account for a significant chunk of electricity consumption in some parts of the US. In six states they consume over 10% of the total power. In Virginia they consume over 25% of total power.

Our own analysis of the 4,500-plus operational and planned data centers in the United States shows hot spots in the cities and regions displayed in the graphic below. Data center developers are attracted to these cities and regions by low energy prices, by supportive local policy, and by available land and water.




aidatacentermap.org

05

# AI's Power Life Cycle



**AI workloads pass through three stages: training, fine-tuning, and inference. Each of these stages consumes electricity differently.**



After preparing the data, the training phase involves feeding the model large amounts of unstructured data. Over time, the AI “learns” the inherent patterns and “learns” to generate the desired output.


Once the model is trained, it undergoes fine-tuning, the phase in which the model is adapted to a specific task. Fine-tuning requires far less energy than the other phases but is “bursty.” That is, it jumps from high to low consumption, causing power spikes.

When the model is ready for use, it is deployed for inference. In this final phase, it receives requests directly from users and applications.

In the case of Google's AI infrastructure, inference represents 60% of total energy use compared to only 40% for training. The high energy use required for inference is due to the billions of requests Google receives for its services.

Other estimates suggest that inference could account for 90% of energy use in Large Language Models (LLMs). [Jegham]

These findings suggest that we can avoid the brunt of the natural resource impact by limiting the use of cloud-based AI to important tasks and avoiding the opt-in of AI-enabled features across platforms.





06

# Estimating Data Center Demands





**Even top researchers admit that calculating data center power and water usage is a conundrum. The speed of the AI industry's growth has allowed data centers to expand faster than transparency and reporting mechanisms can keep up.**



Data centers and their developers treat specifications and technical details of their sites as proprietary information. Estimates of these specs and details range widely, depending on the methodology used. But all academics agree on one thing, data centers are extremely secretive about their power and water usage — rarely sharing public figures, and as we'll see, sometimes not bothering to measure.

In this section, we expand on the components of data center resource usage, exploring different approaches and sharing the formula we have chosen to estimate usage for the AI Data Center Map.

## Efficiency Metrics vs Raw Resource Usage

First, let's look at the efficiency metrics you will typically see in environmental and sustainability reports:

**Power Usage Effectiveness (PUE) and Water Usage Effectiveness (WUE).**


### Power Efficiency vs Power Capacity

The **International Energy Agency** have been tracking PUE year over year.

**PUE** is how much additional power a data center draws on top of the necessary IT equipment power. A perfect score is 1.0, meaning there is no additional energy used for things like auxiliary systems, cooling systems, and offices in the facility.

The IEA data suggests power efficiency is improving in the U.S, however the PUE is a self-reported metric, not measured by independent third parties.

The PUE is typically self-reported by the data centers, not measured by independent third parties, and hide the true scale of power usage.



## Power Capacity

The **power capacity** is total amount of power the data center is rated to use based on its specifications, across these four layers:

- AI chips and core compute (GPUs, TPUs, CPUs)
- supporting IT equipment (networking, storage, memory)
- cooling infrastructure (technologies to pull extreme heat out of the servers)
- auxiliary infrastructure (power delivery, lighting, building management).

When you see that a data center rated at 100 MW, 1 GW, or 10 GW — this figure represents its power capacity. This is the theoretical maximum amount of power it could use, in practice it uses less.

The Lawrence Berkeley National Lab report estimated that AI data centers on average, run at 70% of their operational power capacity. Operational timing also varies, with servers performing training operating 80% of the time, and inference 40% of the operational time.

We have chosen an operational time of 60% as a middle ground, suggesting equal distribution of training and inference.

There is plenty of additional complexity not being accounted for, such as the idle power usage when the servers are not in operation, however we feel these adjustments to a data center's power capacity are sound as a rough estimate.



## Data Centers vs Homes


On the [AI Data Center Map](#), we calculate the number of homes displaced by data center's power usage using this formula:

$$\text{Homes} = \frac{\text{Power Capacity (MW)} \times 1,000 \times 24 \times 0.60 \times 0.75}{\frac{\text{Avg. Home Usage (kWh/yr)}}{365}}$$

We start with the data centers raw power capacity, convert it to kilowatts, and then scale down by the LNBL utilization estimates, to get the daily energy. Divide that by the daily energy usage of a single home to get the equivalent number of homes.

## Try it out!

Visit the [AI Data Center Map](#) use the search bar, or click on a black triangle; a hyperscale data center with a listed power capacity. Click on the info icon next to the Homes Equivalent to see the formula breakdown.



Unknown

Hyperscale

Energy

Status

Company Type

Water Risk Layer

Current Water Stress

Scale

LowMediumHighSevereN/A

**Apple Waukeet Campus**

Developer: Apple

Status: Operational

Operating Since: 2024

Power: 200 MW

Est. Power/Day: 2,016,000 kWh

Est. Water/Day: 239,657 gal

Est. Homes Displaced: 70,080

Est. Heat Island: +3.6°F

Type: Hyperscale

Location: Dallas County, IA

Sources: [1] [2] [3] [4] [5] [6] [7] [8] [9]

### Est. Homes Displaced

Step-by-step formula breakdown:

|                              |                      |
|------------------------------|----------------------|
| Power capacity               | 200 MW               |
| × 1,000 kW/MW                | 200,000 kW           |
| × 24 hrs/day                 | 4,800,000 kWh/day    |
| × 70% op. power [LBNL, 2024] | 3,360,000 kWh        |
| × 60% op. time [LBNL, 2024]  | 2,016,000 kWh        |
| <b>Est. Power/Day</b>        | <b>2,016,000 kWh</b> |
| Avg home usage [EIA, 2024]   | 10,500 kWh/yr        |
| ÷ 365 days                   | ~28.77 kWh/day       |
| 2,016,000 ÷ 28.77            | 70,080 homes         |
| <b>Est. Homes Displaced</b>  | <b>70,080</b>        |

Adjusted for avg. server operational power (70%) and time (60%). Home usage based on U.S. avg of 10,500 kWh/year.

## Water Usage Effectiveness (WUE)


Hyperscale data centers are special in their extreme water use. Large amounts of water is not only needed to cool the servers — but also at the power plants which produce the power.

**Water Usage Effectiveness (WUE)** is the amount of water needed per unit of power, expressed typically in liters per kilowatt-hour.

A 1 GW data center with an WUE of 0.3 L/kWh means approximately 7.2 million gallons of water use per day (not factoring in operational time or power).

And the water usage on site is only one half of the total water impact, in reality there are two **WUE ratios** to discuss.

- **Direct (On-Site) WUE** — Measures how much water is needed for on-site cooling
- **Source (Upstream) WUE** — The water used to produce the energy for the data center (or the inputs for manufacturing energy inputs)



The Source WUE can vary wildly due to the mixed supply of power sources for a given hyperscale data center. Each power source, whether it be natural gas plants, nuclear, or geothermal has its own levels of water demands.

It is the metric rarely mentioned in conversation, and absent from transparency reports which only share the Direct On-Site WUE. We must acknowledge both halves of WUE, some researchers suggest that the **Source WUE**, is the smaller part of the impact.

Additionally both **WUE metrics** will range widely on a multitude of factors including the type of cooling technology, the location of the hyperscaler, the time of year, and the surrounding environment.

Water usage is highest in hotter, humid areas, in the hottest seasons.

In addition, **Source WUE** is a self-reported metric, published in the specification or sustainability report. In many cases it may be a simple estimate, in the next section we'll cover how many data centers aren't tracking their usage.

## Try it out!

Visit the [AI Data Center Map](#) and click on a black triangle; click the question mark next to Est. Water/Day to get an estimate of its water use.

## Which Estimate is Correct?

There are many papers attempting to estimate the power and water usage of data centers, however we (and most academic researchers) have chosen the Lawrence Berkeley National Lab approach due to their creative bottom-up modeling, which circumvents the lack of power and water use transparency by data center operators.





## Here's an overview of their process:

1. **Collect** all power specifications for data-center components (IT equipment, AI chips, storage, networking, cooling, and auxiliary systems).
2. **Cross-reference** market import data to determine the quantities of equipment purchased and brought into the U.S.
3. **Multiply** the specification-based power usage by the national install base to obtain rough total power consumption.
4. **Apply** thermodynamic modeling to estimate the water required for different cooling systems and climate zones.

07

# Big Tech's Sustainability Facade



**Big tech companies and data center operators will be quick to set the environmental narrative in their yearly sustainability reports. Its common to see these reports dressed-up with uplifting visuals of newly created jobs, clean energy, social programs, ecological projects (that never seem to be in the same areas as the data centers), and finally efficiency metrics like PUE and WUE that hide the whole truth.**

## **Data Centers May Not Be Measuring Water Usage**

Below is Google's 2025 Environmental Report, issued in June 2025 and containing 2024 data. It shows that the company used 11 billion gallons of water in 2024. Even this staggering number isn't an accurate measurement; it's a self-reported prediction.



aidatacentermap.org

| Metric<br>Million Gallons | 2020  | 2021  | 2022  | 2023  | 2024<br>Total | 2024<br>Offices | 2024<br>Data Centers |
|---------------------------|-------|-------|-------|-------|---------------|-----------------|----------------------|
| Water<br>withdrawal       | 5,689 | 6,297 | 7,600 | 8,653 | 11,011        | 1,145           | 9,866                |

\*Water consumption = water not returned to the source.

Reading the fine print of the appendix, we find that water usage across Google's facilities is sometimes estimated. In fact, sometimes it's not even being measured.

We calculate **water consumption** by subtracting water discharge from water withdrawal.

**Water withdrawal** is based on actual metered or invoiced data when it's available. At offices where actual metered or invoiced data isn't available, we estimate water withdrawal using facility square footage and internally developed water withdrawal intensity factors by office type based on data from the reporting period. At data centers where actual data isn't available, we estimate water withdrawal using engineering principles.

**Water discharge** is based on actual metered or invoiced data when it's available. Where actual domestic wastewater discharge data isn't available, we apply an industry-standard 90% discharge flow factor to a facility's domestic water withdrawal to estimate domestic water discharge and a 0% discharge flow factor to a facility's irrigation water withdrawal to estimate irrigation water discharge. We apply this water discharge estimation methodology at all offices and at data centers where metered water discharge data isn't available. At data centers where actual data isn't available, we estimate water discharge using engineering principles.

aidatacentermap.org

Buried in the Appendix of Google's 2025 Environmental Report

Instead, Google estimates water withdrawal using square footage, internal factors, and "engineering principles." In other words, it's a black box estimate closed off to third-party validation.

## Lawrence Berkeley National Lab estimates that less than one third of data center owners in the US measure their water consumption.

A separate 2021 study by Uptime found that 50% of data centers surveyed were tracking water use in some way. Less than 10% tracked water usage across their entire fleet.

Why the hesitation? “Most say it’s because there is no business justification, which suggests a low priority for management.” [Uptime, 2021]

With billions in their capital budgets, these data center operators have the means to install basic water-metering. There’s no excuse for the avoidance of measuring the actual numbers—which if exposed may reveal the truth that their data centers are draining communities of precious water resources.

Considering that two-thirds of data centers in the United States are built in areas suffering from water-stress, these companies have an undeniable obligation to measure and disclose their use of water, our most precious resource on earth.



## Sustainability Report Breakdown

When we look at transparency/sustainability reports in this industry, we find they leave out accurate information, hiding behind estimations and efficiency metrics. We've looked at the sustainability reports published in the year 2025, detailing the resource usage for the previous year (2024).

**Google** — Of all the big tech players, Google is among the most transparent. They publish their total power use (32 TWh), total water withdrawal (11 billion gallons), and power and water use by site. It's unclear how many facilities estimate usage versus report true measurements.

**Meta** — Meta's report showcases a long list of carbon emissions, renewables, and biodiversity programs. You must look at their **environmental data index** to see 18

TWh of power usage and 5.625 billion liters of water use total. Power and water usage per site is broken down, but it's unclear how many facilities estimate usage versus report true measurements.

**Microsoft** — Microsoft publishes totals: a global 30.207 TWh of power and 10.377 billion liters of water. Their water comes mainly from third-party water sources (i.e., local municipalities).

**Amazon** — Amazon publishes a sustainability report but provides neither power nor water usage figures. They only provide efficiency metrics rather than revealing per-facility breakdowns.

**X (Twitter)** — After going private in 2022, X was no longer obligated to publish environmental disclosures. We were unable to find any transparency or environmental reporting in recent years.


**Anthropic and OpenAI** — Historically these two AI giants rented compute capacity from cloud providers. As of 2026, both are building data center capacity of their own. We have been unable to find any formal sustainability or environmental reporting from either company at the time of writing."

# Who's actually measuring their environmental impact?


Tech companies rarely disclose power and water usage, and when they do, they may not even be measuring.

## Partial Disclosure: Totals Published, Metering Details Missing.

### ⚡ Power Use 2024



### 💧 Water Withdrawal 2024



Totals published ≠ fully measured. Metering coverage is not disclosed.  
Water withdrawal may include estimates rather than direct measurement.

## No Meaningful Disclosure

Amazon

EFFICIENCY METRICS PUBLISHED, MISSING POWER,  
WATER TOTALS AND PER FACILITY BREAKDOWN

x (twitter)

NO ENVIRONMENTAL REPORTING FOUND.

Anthropic

NO ENVIRONMENTAL REPORTING FOUND.

OpenAI

NO ENVIRONMENTAL REPORTING FOUND.

## Comparing Water Use

Some operators frame their water using relatable comparisons. For instance, Google uses the term “golf course equivalents” to describe the water impact of each data center.

Google frames data center water use through a high-water benchmark that makes millions of gallons seem smaller. The table’s visual cues, including the check marks, further reinforce that framing.

### CASE STUDY COUNCIL BLUFFS, IA • 2024

**1,010.2** million gallons → **6.7** golf course equivalents

based on an estimated ~150M gallons per golf course per year

| Water use by data center location |                 | 2024                      |   |           |   |             |   | Golf course equivalents (estimated) <sup>165</sup> |
|-----------------------------------|-----------------|---------------------------|---|-----------|---|-------------|---|----------------------------------------------------|
| Location                          | Unit            | Withdrawal <sup>164</sup> |   | Discharge |   | Consumption |   |                                                    |
| Ashburn, VA                       | Million gallons | 59.5                      | ✓ | 3.5       | ✓ | 56.0        | ✓ | <1                                                 |
| Berkeley County, SC               | Million gallons | 853.8                     | ✓ | 77.3      | ✓ | 776.5       | ✓ | 5.2                                                |
| Bristow, VA                       | Million gallons | 105.7                     | ✓ | 21.3      | ✓ | 84.4        | ✓ | <1                                                 |
| Council Bluffs, IA                | Million gallons | 1,410.3                   | ✓ | 400.1     | ✓ | 1,010.2     | ✓ | 6.7                                                |
| The Dalles, OR                    | Million gallons | 461.1                     | ✓ | 99.7      | ✓ | 361.4       | ✓ | 2.4                                                |
| Douglas County, GA                | Million gallons | 444.1                     | ✓ | 77.2      | ✓ | 366.9       | ✓ | 2.4                                                |

**Context matters.** Southwestern golf courses are highly water-intensive and often exceed water allotments.

[aidatacentermap.org](https://aidatacentermap.org)

**In our view, comparisons like this distract from the raw water use. A reader may think, “how much water could a golf course actually use?”**

Reading the fine print, we find that Google is referring to large golf-courses in the arid Southwest, which require the most water than anywhere in the country, a whopping 150 million gallons of fresh water annually for an average golf course.

This comparison doesn't make data centers look better, it makes golf courses look worse.

And there is no surprise that golf courses in the American Southwest are facing scrutiny. The **Arizona Republic found** that 30–50% of Arizona golf courses have exceeded their water allotments in any given year for the past 20 years. 86 golf courses have exceeded allotments at least 10 times since 2002.

Despite these golf courses claiming themselves sustainable (sound familiar?), the findings showed they made no collective progress in reducing water consumption in 20 years and continue to exceed allotments without consequence.



## Conclusion

The organizations meant to responsibly manage the water have little ability to enforce rules. If local water authorities can't keep golf courses in check, how could we expect them to keep billion dollar hyperscale projects in check?

There is a massive accountability gap for water abusers, and its unlikely we can count on local water development boards to stop abuse.

## Section Recap

### Can transparency reports be trusted?

We find that major data center developers cannot be trusted to share accurate, complete information:

- There is no clarity on which data centers are metered and which ones are estimated
- They hide behind efficiency metrics and avoid sharing raw usage
- Their sustainability reports focus on water conservation projects without clarifying how water tables recharge themselves with precipitation and with water leaking from rivers, lakes, and ponds
- They contrast themselves with worse water-guzzling industries, like golf courses, which similarly abuse their water allotments without consequences and which haven't made any serious progress in reducing water usage

08

# Our Water Systems






# Its time to turn our attention back to nature, let's review how our natural water systems actually work.

If you go to the beach and dig a hole in the sand, you'll hit water almost immediately. This is really no different then tapping into an aquifer, or bodies of water located underneath the ground, constrained by impermeable rock.

A single aquifer can span thousands of miles underneath the ground, spanning multiple states or even countries. Or they can be smaller than a square mile, supplying a few properties with well water.



Inside our earth, different types of rocks have different abilities to store and carry water, depending on their permeability. Water travels through permeable rock, making its through underground rock formations, eventually pooling at a stopping point called the “zone of saturation”, or the store of water for an aquifer that wells and springs connect to.

In this zone, the top-level surface surface of the water is called the water table, measured in the number of feet below the surface it sits.

Wells require drilling a path down to the aquifer, and equipment to pump the water upwards.

Springs are open paths to the surface where water naturally flows upwards, due to the weight of the earth on the lower aquifer.

## How Aquifers Grow and Deplete

Our aquifer and groundwater can be thought of like a bank account. The amount of water stored can be measured by the water table is our current savings balance.



This balance is remains relatively stable, but as new sources of demand come in, such as heavy industry or the millions of gallons required by hyperscale data centers — a big bill needs to be paid.

As water-consumers pull from wells, the water table drops across the aquifer. If the water table drops significantly, it could cause wells to run dry and damage to the ecosystem.

How do we restore the water supply?


Aquifers recharge when water is returned to recharge zones, by way of rain or when water is manually diverted to these areas. These recharge zones have permeable rock on the surface, providing an unimpeded route for water to travel down into the aquifer.

On the way down, the water may travel through different aquifer layers like sand and porous rocks, filtering the water of toxins and pollutants.

Water kickstarts the cycle of life, and hyperscale data centers may put everything at risk.

## The Edwards Aquifer Region

### San Antonio and Barton Springs Segments



## Groundwater Conservation

The Edwards Aquifer is one of the largest artesian aquifers in the world, a primary source of drinking water in Central Texas. We'll look at the Edwards Aquifer to understand the common zones of an aquifer.

The Edwards Aquifer is a fault zone aquifer, however there are other types. Each aquifer's specific type and geology should be carefully studied.

The contributing zone (or catchment area) “catches” rainwater, which starts its journey by entering rivers, streams, or creeks.

These waterways carry the water into the recharge zone, where the water disappears into the ground, making its way through permeable rock as it sinks downwards into the aquifer. The Edwards Aquifer recharge zone has highly fractured limestone on the surface.


Water conservation can involve routing water to these recharge zones, and it's critical to maintaining the health of the aquifer.

For example, Valdina Farms Sinkhole in Texas receives heavy rains from nearby creeks, recharging the Edwards Aquifer at 1,770 gallons per second.



Valdina Farms Sinkhole

# Typical Cross-Section of The Edwards Aquifer



The water finds its way down into fractured limestone and eventually becomes “trapped” inside. The weight of the rocks and the incoming water from the recharge zone combine to apply a huge amount of pressure on water in the deeper areas.

This area of water under pressure is called the artesian zone. In the artesian zone, the surrounding pressure pushes the water up through cavities or through drilled wells.



In the late-19th century, the first wells drilled into the Edwards aquifer released a massive amount of pressure, spouting water 25 feet or higher into the air. Over decades of water consumption, the pressure has faded and requires the use of manual pumps to draw what is left.

Now that you have a basic understanding of aquifers, you can view the “U.S. Major Aquifers” layer on the AI Data Center Map to see how many data centers are positioned over—or next to—aquifer zones.



1891 San Antonio, TX



## Restoring Our Aquifers

**We should point out that even though aquifers recharge, this recharging process happens very slowly. It can take hundreds of years of responsible water use for an aquifer to return to normal levels.**

The water in aquifers is ancient. It has been stored underground for thousands of years. When hyperscale data centers slurp up millions of gallons a day from aquifers, its unable to recharge. The water is passed through extremely high temperature cooling environments and released into the atmosphere as steam.

Its impossible for aquifers to recharge at the same rates as hyperscale water consumption, it will take decades of rain and water conservation to restore ecological balance.

We'll see in the next section how the US and the rest of the world have been dealing with chronic groundwater depletion.

09


—

# Dangers To Water Systems



The United States has been facing severe groundwater decline since the 20th century as a result of its rapid development and use in agriculture, heavy industry, and now data centers.

Over half of the United States relies on groundwater for drinking water. If we continue at our current unchecked pace of development, we may begin to see the collapse of natural water systems.



1988 - 1997, declines of 300 feet in Santa Fe, NM

**Water levels can drop drastically within a few decades of intensive water use. Studies use the depth of wells to track water depletion over time. Since the 1980s, different areas of the U.S have experienced drops in the groundwater level anywhere from 100 feet (in the High Plains) to 900 feet (in the Great Lakes area).**

Heavy industries such as power generation, electric vehicle production, agriculture, are consuming similar amounts of water to hyperscale data centers, in total pumping tens of millions of gallons a day, leading to an extreme water deficit.

## **Impacts Of Severe Ground Water Depletion**

Over the years, the United States Geological Service (USGS) has extensively studied groundwater trends and impacts resulting from groundwater depletion.

A 2024 USGS study of global groundwater trends across 1,693 aquifer systems worldwide concluded that the



groundwater level is declining rapidly—up to 1.5 feet every year in many regions.

Despite this worrisome trend, the analysis pointed out that groundwater decline can be reversed with proper management and routing of surface water. There is hope, but only if we put a stop to the currently unchecked growth of data centers.

## Try it out:

Click the dropdown on the top left to visualize aquifer layers on [AI Data Center Map](#). What's your nearest aquifer and what hyperscale data centers are located over it?

## Water Wells

Property owners in areas with hyperscale data centers are going to see a real and lasting impact on their water wells. As groundwater is depleted, it can fall below the level of currently dug wells, water output will slow, lessen, and even completely dry up.

A dry well must be re-drilled and deepened, this is an expensive operation that could potentially ruin a farmer or property owner who's dependent on their well.



## Try it out:

Use the [AI Data Center Map](#) to estimate the amount of water displacement near you by clicking on an empty space on the map.

- How many home equivalents are being drained?
- What hyperscalers are in your area?

## Reduced Surface Water

Our lakes, rivers, and wetlands rely on groundwater sufficient to fill and recharge them. Declining groundwater levels means that less of it flows into these bodies of water.


We see the results of groundwater depletion on the earth's surface as smaller lakes, thinning rivers, and fading ponds. These effects may take many years to develop.

## Environmental Loss

When less water is available above and below ground, the environment suffers. One area of environmental focus takes place in what is termed the riparian zone. The riparian zone is the land directly adjacent to a river or stream, the special transition zone between the body of water and the land.

The riparian zone is where land meets water, figuratively and literally holding our environment together.

In this zone water-loving vegetation forms a powerful network of roots that prevents soil from eroding. New soil is generated here as well, the moist conditions leading to the quick decomposition of plant matter combining with sediments and processed by bacteria, fungi, and worms.



Tuscon, Arizona, riparian trees have  
disappeared after 47 years

Pollutants and toxins are also filtered out from the thick plant mass, before they enter our drinking water.

Trees next to the riverbank provide shade for the families of fish and amphibians who live in this zone, and provide a respite for animals and water birds who live in the region or are passing by on their migratory routes.

**When hyperscale data centers deplete water levels, the waters recede; putting plants, animals, and ultimately people at risk.**

The most severe groundwater depletion are in arid or semi-arid climates with high amounts of cultivated land. These areas take even longer than normal to recharge and recover from depletion.

## **Subsidence**

When water levels decline, the earth feels the loss of its structural contribution and can undergo shifts underneath the surface, causing the earth to settle or sink. It is estimated that 80% of subsidence in the US is related to groundwater depletion.

## Toxins & Deteriorating Water Quality

Both seaward and inland aquifers can suffer from extreme pumping. Over-pumping displaces the water and can create two major problems for drinkable groundwater. As large amounts of water are pulled upwards:

- Layers of salinated water below the aquifer can be pulled upward and into the drinking water supply
- Polluted water can be pulled downwards from the surface nearby, creating a “drawdown” effect

Water can become polluted when it is drawn down—such as when heavy industry builds a facility on top of an aquifer’s recharge zone and pollutes the lands surface. Rain mixes with surface pollution, and heavy pumping in the area can pump this down into the aquifer, circulating into the drinking water.

This typically occurs when confining layers are damaged, or layers of earth meant to prevent cross-contamination of aquifers. These layers “tear,” thereby allowing polluted water to move through the holes and mix in with the main source of drinking water.

**ProtectOurAquifer Memphis** constructed a map to illustrate suspected confining layer breakdowns in that Tennessee metropolis. Tears in the area can be identified by testing water and identifying artifacts that should not be present in natural groundwater that is left undisturbed.







10


U.S.

Water Rights



Groundwater is a shared resource, and aquifers are shared across property lines, state lines, and even country lines. Those with the rights to pump water can abuse them, reducing water levels and impacting the water accessibility for their neighbors.

In the U.S, each state uses one of four groundwater governance systems. A 2020 joint legal “survey” from Texas A&M University and the National Agricultural Law Center’s provided an overview of these four systems and where they are used.



# Four Types Of Groundwater Governance

## 1. Rule of Capture

Rule of Capture, also known as “absolute dominion,” is a doctrine that gives property owners complete ownership of the water under their lands.

Texas is the only state that uses an unmodified absolute dominion system.



**According to Gabriel Eckstein, J.D., a professor of law at Texas A&M University School of Law:**

**The Texas approach is: if you own the land, then you own the groundwater that's underneath your feet.**

Many states started off with this rule but evolved over time as water became more scarce.

This means that data centers in Texas can use as much water as they want, potentially ruining the water rights of others without facing any legal or economic consequences.

Other states that apply the absolute dominion doctrine, but to a lesser degree, include Indiana, Louisiana, Connecticut, and Massachusetts. Each of these states has made various modifications to the rule, such as restrictions, permits, and registration requirements.

Its clear why hyperscale data centers are so keen to develop in these states.

## 2. Reasonable Use

The Reasonable Use doctrine, used by seventeen states, is similar to absolute dominion, but water rights-holders may not interfere with another's water usage or unreasonably harm neighboring wells or aquifers.

## 3. Prior Appropriation

Under the Prior Appropriation doctrine, which is governed by the principle of "first in time, first in rights," water rights-holders are prioritized based on when they started using the water for a beneficial use.

Senior claimants are preferred to junior in times of water scarcity. Senior claimants receive 100% of their

their water stake, and after all of the water budget is accounted for, junior claimants receive 0% of their stake, effectively unable to pump water.

In many states this doctrine has been replaced by a permit system. Thirteen states have adopted or favor Prior Appropriation today.

## 4. Correlative Rights

Combining the previous three approaches, the Correlative Rights doctrine gives groundwater rights-holders a proportional share during scarcity. The landowners directly over the aquifer are given priority over those not located on the aquifer.

# Texas And Groundwater Conservation Districts

Texas has adhered to the absolute dominion doctrine in groundwater lawsuits that have reached the Texas Supreme Court. In Texas, if you can pump it from a well on your land, it's your water.





**Amy Hardberger, professor of water law at Texas Tech University School of Law, observes:**

Texans need to understand the rule of capture and that, generally, [it] allows a tremendous amount of pumping legally, which can cause problems. [...] And then, [they need to] understand that there might be regional regulation that may change that. If every Texan even understood that very, very basic thing, we would be better off. Unfortunately, a lot of times, people don't realize how we regulate groundwater until it hurts them.

Project Stargate data center in Abilene, Texas, or Project Matador in the Texas Panhandle show us the flaws in the absolute dominion model.

These two data centers sit on the Ogallala's aquifer, one of the largest aquifers on earth serving eight states across 170,000+ square miles.

The fact that data centers in Texas can pump as much groundwater as they want without any consequences or accountability, threatens to gravely harm millions of people relying on groundwater.

## Groundwater Conservation Districts

Also unique to Texas are its Groundwater Conservation Districts (GCDs). These districts enable a local board of directors to manage and protect groundwater across the state.

Every five years, the ninety-eight Groundwater Conservation Districts throughout the Lone Star state come together to plan water resources for the next fifty years, aiming to establish and update the “Desired Future Conditions” for relevant aquifers.

With these conditions in mind, they assess current water demands and set guidelines for new groundwater permits and allowable withdrawals for the upcoming period.

However, there’s a major problem. The last Desired Future Conditions meeting took place in 2021, and did not account for any data centers in the State of Texas. Moreover, GCD representatives won’t meet again until 2027 to publish their next report.

In that time, there’s been an extremely rapid buildout of data centers.

Although GCDs monitor groundwater usage and production, they have very little information on the private power- and water-contract arrangements between utilities or municipalities and the data centers that are already operating or planned—contracts that are typically under NDA.

Without this information, they risk over-allocating water resources, leaving groundwater issues in a gray area until the next meeting. Until then, there is no official guidance on the impact of data centers.

## Other States

**Other states also have organizations that play a role similar to Texas' Groundwater Conservation Districts. For example:**

- California has Groundwater Sustainability Agencies (GSAs).
- Arizona, New Mexico, and Colorado have state engineer offices and designated basins.
- Eastern states use state environmental agencies and water permitting programs.

11

# Hyperscale Life Cycles: Community Wins And Losses





In the spring of 2026, residents of San Marcos, Texas, came out for the second time to **reject 904 Francis Harris Lane, a \$1.5-billion planned data center project.**



The 200-acre site of the project would span five buildings and use 380 megawatts of power—almost two-and-a-half times the power of the entire city of San Marcos at peak demand.



**We believe this is one of the most important case studies on data centers in the United States. It reflects community-wide efforts to protect the land that sits atop the Edwards Aquifer, home to crucial Texas Hill Country water systems such as the Edwards Aquifer, San Marcos Springs, the Blanco River, and the Guadalupe River.**

We'll use this case study to explain how a data center goes from idea to completion—and how concerned individuals can stop it if they step in at the right time.



## **Real World Examples**

Look out for the concise info-boxes that anchor each stage of the lifecycle with a real world case study.

# 1. Site Selection


A hyperscale data center must require access to power, and a lot of it. They connect directly to the high-voltage backbone of the electric grid; a 345 kilovolt (kV) transmission line.

Hyperscales are built within short distances of these lines, extending high-voltage lines is expensive (\$1M per mile) and permitting can take years.

Redundant power is crucial, the ideal location would be adjacent to a gas, hydroelectric, or geo-thermal plant with extra capacity.

**Armed with this information, you can predict where data centers are likely to spring up.**

Dallas data centers hugging high voltage transmission lines





## Pre-2025 — Site Selection: 904 Francis Harris Lane, San Marcos, Texas

904 Francis Harris Lane was converted from a failed neighborhood development into a data center. It benefited from being adjacent to the Hays Power Plant and high-voltage transmission lines.

## 2. Going Industrial

Once the land is purchased, the developer may need to rezone the land from residential or agricultural.

Industrial zoning enables property owners to build on-site power generation (e.g. gas turbines), larger water mains, and larger building footprints without having to jump through restrictions or permits.

Industrial land is typically not available at the scale required for data centers, which require hundreds if not thousands of acres. For this reason, it's common for data centers to rezone from other types of land.



## Early 2025 — Initial Highlander SM One LLC Proposal

Data center developer Highlander SM One LLC requested San Marcos City Council to:

- Annex 64 acres from the county into the city of San Marcos
- Amend the San Marcos desired map to turn the site from conservation to commercial
- Most importantly, rezone 200 acres from “Future Development” to “Light Industrial”

## Planning And Zoning Commissions

Rezoning is decided at **planning and zoning commission hearings**. This step is critical because its the only one that **requires a public hearing**, where residents can speak on the **official record**.



## March 2025 — San Marcos Planning and Zoning Commission Hearing

The San Marcos Planning and Zoning Commission recommended a denial for all of Highlander SM One's requests, including rezoning. The rezoning application then moved on to the San Marcos City Council.

However the zoning commission recommendation is not the final say, it is advisory. Depending on local processes their decision may influence the amount of votes needed in the final voting body.





## August 2025 — San Marcos City Council Hearing

In San Marcos' case, since the Planning and Zoning Commission denied the application, the City Council had to approve it with a supermajority (that is, 6-out-of-7 members voting "yes").

### This hearing resulted in:

- Unanimous approval for annexation, bringing the site within the city's jurisdiction
- A 5-to-2 City Council vote in favor of the map change, but ultimately failing to reach a supermajority
- Because the map change failed, the City Council could not vote on rezoning.

**With this rejection, the application could not be voted on again for another six months.**

## Jurisdiction

Notable here is how the San Marcos City Council annexed the data center land, which brought it fully under the city's jurisdiction.

### Data centers may fall under different jurisdictions in the United States:

- **Within city limits**

Here The data center falls under city authority: the city's planning and zoning commission hears any applications, and the city council makes the final decision.

- **Within Extraterritorial Jurisdiction (ETJ)**

Cities of a certain size are given planning and zoning authority beyond their city limits. ETJ is handled differently in different states.

- **Outside city limits**

County land or unincorporated land falls under the county government's jurisdiction. The county's planning and zoning commission (if it exists) hears the application, and the county commissioners decide.

Landowners in a county or an unincorporated area can petition a city to annex their land. In this case, the data center developer Highlander SM One got a twofold benefit:

- Access to San Marcos city water and the city's sewer system
- A simplified approval process in a single jurisdiction (the city)



## October 2025 — New Application

Highlander SM One LLC did not bother to wait the full six months. They instead filed a new application for the exact same project, minus the previous operator. The project was scheduled for zoning review in a January 2026 public hearing with the two remaining requests (map change, rezoning).

**In a strange turn of events, the same zoning commission that rejected the rezoning the first time, switched their vote.**



## January 2026 — Rezoning Approved

In this second Planning and Zoning Commission meeting, members gave a positive recommendation to rezone the property for data center use. The rezoning request moved on to the San Marcos City Council. This time, due to Planning and Zoning's positive recommendation, it required only a simple majority (not a supermajority) vote from the City Council.

At this point, all that stood between this town and a new massive data center was a single city council vote.

Multiple local organizations sprang into action and began raising local awareness. A petition with more than 4,000 signatures was created, protests were held, and the town was abuzz with news of the upcoming hearing.



## February 2026 — City Council Denied Rezoning

The final San Marcos City Council rezoning meeting lasted eight hours. More than 130 public comments were made at the public hearing, which ended at 1:00 AM. Council members issued their final decision after 3:00 AM. The rezoning was denied by a vote of 5 to 2. Highlander SM One would not be able to refile its application until August 2026.

## Is It Over?

This last vote by the San Marcos City Council was a fairly decisive victory for data center opponents. Even if the developer, Highlander SM One, makes its way through the Planning and Zoning Commission, it will still need to change three votes on the City Council to win.

Still, locals shouldn't let their guard down.



Although it's looking like an uphill battle for the developer, if new members are voted onto City Council, it is possible the project could be approved.

But, unlike San Marcos, not all communities are blessed with a data center that can be meaningfully resisted.

### 3. Tax Breaks

CyrusOne's DFW10 data center, roughly 25 miles north of Waco, is built on unincorporated land and is outside city limits. Thus, there is nothing Bosque county can do to prevent the data center from being built. Instead, the county chose to give a ten-year 30% property tax abatement, which the county can use as "leverage" against the data center operator.

**There are two major areas where data centers can get tax breaks:**

#### 1. State-level sales and use tax exemptions:

This exemption provides a tax break on equipment going into the data center, saving large hyperscale projects billions of dollars.

#### 2. Local property tax abatements:

As is the case with the Bosque County data center, a county can excuse a proportion of property taxes.



## Texas' Chapter 312 Property Tax Abatement

According to the Texas Comptroller's Office, Chapter 312 tax abatements are an economic development tool to reduce unemployment and increase tax revenue for the area.

However the deals put together by cities and counties often have no hard requirements and end up passing additional costs back to taxpayers.

We were able to find the details of the CyrusOne data center by looking at the minutes of the Bosque County Commissioners Court.

## Bosque County Tax Abatement Details

- Waives 30% of the county property taxes for ten years
- Applies to new buildings, structures, site improvements, offices but does not apply to equipment like computer servers, inventory, and supplies
- Makes the operator (CyrusOne) responsible for all damage to roads from construction
- Allows the county to inspect the site with 48 hours notice
- Requests the data center to make reasonable efforts to use local vendors and businesses, though this is not a hard-and-fast mandate
- **There is no stringent requirement to hire locally.**

**CyrusOne commented that the deal with Bosque County was the least impressive out of their many projects.**

## Controversial Tax Abatements

There are far more controversial deals than Bosque County. In Louisiana, Meta received a 100% exemption of state and local sales taxes for most data center costs including AI chips, servers, and cooling infrastructure.

## 4. Power And Energy Contracts

Purchasing a property in close proximity to high voltage power lines does not guarantee access to power. Data centers must secure their power and energy contracts with an Interconnection Agreement, which connects them to the electric grid.

Organizations called Independent System Operators (ISOs) or Regional Transmission Organizations (RTO) manage and maintain their grids, selling capacity, and balancing supply and demand. As new power requests come in, these organizations perform power studies to ensure power stays reliable.

From 2022 onward, multiple regional grids have become overwhelmed with new requests. They have openly admitted that their approval processes have broken down, some requests have taken up to eight years to be granted (or denied) — including new power sources.

**For example, at its December 2025 auction for the 2027–2028 delivery year, PJM Interconnection, a regional transmission organization that manages the largest power grid in the US, came up 6.6 GW short of its 20% installed reserve margin target.**

The target enables the electric grid to remain stable by having 20% more electricity generation than forecasted peak demand.

The culprit? “Phantom load” is mainly to blame. Data center developers are playing the field, keeping their options open by applying for multiple grid interconnections across multiple electric grids, even if they’re not intending to build.

This has been causing a massive inflation in forecasting and wholesale electricity prices.

As a result, grid operators have begun implementing emergency rules that hold data center developers accountable for their power requests and ensure they commit financially up front by contributing to the costs of interconnection and transmission.



## These new rules also enable creative new strategies for data centers to become part of the grid faster through three new options:

### Co-Location:

Data centers are built on-site or right next to a power generation station (a natural gas, hydroelectric, solar, or nuclear power plant) so they can use that plant as their primary power source and reduce their demand requests on the main grid. This comes with stipulations depending on the grid. For example, in Texas co-located data centers must shut down their operations within 30 minutes of a grid emergency.

### Power Purchase Agreements (PPAs):

Under these long-term contracts with independent power producers, hyperscalers are getting creative by securing renewable energy and nuclear power, relying less on grid power. However, stipulations on some grids (Texas' ERCOT) prevent hyperscalers from using any remote power during emergencies.

### Off-Grid:

To avoid the interconnection process entirely, some data centers are now powering their facilities completely off the grid. This is typically done by installing large numbers of natural gas turbines.

## Project Goodnight Co-Location

East of Amarillo, Texas, is Project Goodnight, a \$29-billion hyperscale data center being developed by Crusoe Energy (of Claude, Texas) in collaboration with Google.

The project's seven-building Goodnight Campus, which aims to provide 1GW of capacity intended to support AI infrastructure, is described as a "copy-exact twin" of Crusoe's nearby data center in Abilene, Texas, which is tied to OpenAI's Project Stargate.

Project Goodnight uses two different power supply avenues that are under different regulations.

### 1. Off-Grid:

Crusoe plans to build a natural gas plant right next to the campus to supply two buildings with power. These gas turbines will be off-grid, not connected to the Electric Reliability Council of Texas (ERCOT), so they don't need to go through an interconnection process.

### 1. Co-location:

Another part of the campus will be powered in part by Serena Energy's Goodnight Wind Farm, which is connected to the grid and is the first project of its kind to deal with ERCOT's newly passed Texas Senate Bill 6 (SB6) regulation.

**SB 6 mandates that in times of grid emergency, the data center has 30 minutes to shut down its power draw so the wind farm can contribute back to the grid. In the case of such an emergency, Crusoe will not be paid for the power it provides back to the grid. Additionally, the data center cannot leverage off-site power plants to continue operations. Instead, it must shut down.**

**Instead, it must shut down.**

Despite regulatory bodies making it harder for data center developers to make phantom power requests, the damage has been done. Demand has been forecast, energy costs have skyrocketed, and consumers will pay billions as a result.

## 5. Environmental Permits

In the US, data center developers will request air quality and water permits from state environmental agencies or county agencies.

In addition to grid power or on-site power generation, a hyperscale data center can have anywhere from 15 to 700 (or even more) generators on site, used for auxiliary and backup power

Each one of these generators is a mini power plant and is powered by diesel fuel or liquid natural gas (LNG).

In the past, data centers used mobile power plants only for backup power. These days, though, companies like xAI have found creative ways to permit generators temporarily while building out a standard power plant.





## March–June 2024 — Powering xAI Colossus in Memphis

In mid 2024, the xAI Colossus site was confirmed in the South Memphis part of Memphis, Tennessee. But the grid connection in this area was only 8 MW—nowhere close to the power needs of the site. It was reported that xAI brought in over a dozen portable natural gas turbines, and operated them without any air permits.

Developers will apply for a permit from their state environmental agency.

**There are two major classes of air permits, according to the US Clean Air Act:**

### Major Source Permits:

Industries that produce more than 100 tons per year (TPY) of air pollutants in 28 specific categories, or more than 250 TPY of any pollutant, are categorized as a “major source.” These “major source” permits have demanding regulations, including pre-construction permits, evaluation of all pollutants and emissions, and air quality impact analyses.



### Minor Source Permits:

Any industry producing less than the cutoffs of Major Source Permits can become a minor source. They do not have to undergo pre-construction permits and evaluations.

Data center developers want to avoid being classified as a major source due to the regulatory demands.



## January 2025 — xAI applied for Minor Source Permits for Colossus 1

After announcing its intention to become the largest super computer in the world, xAI's grid allocation is raised to 150 MW, with a large fleet of turbines to power the rest of the demand. Despite all the turbines, xAI applies for a minor-source permit(s) for the facility.

The minor permit was later approved by Shelby County Health Department for 15 turbines, a **decision appealed by the NAACP and SELC.**

## Data center operators like xAI can utilize loopholes to obtain a minor source permit while having a fleet of gas turbines.

Under the Clean Air Act, source classification is based on tons of pollutions emitted per year. However, regulators use the generator running hours as a reportable metric to calculate the emissions are under the permitted thresholds.

### Let's look at the loopholes:

1. Permit each turbine individually instead of a fleet, so the cumulative impact is not accounted for. xAI did this with each of their turbines, ensuring each turbine appeared to be a minor source.
2. Operators are allowed to define emergencies, and operation during emergencies is not counted towards operating hours. Operators can declare emergencies such as grid voltage fluctuation, grid maintenance, maximum generation, severe weather, and increased demand — and not have this count towards emissions.
3. Treat their turbines as 'temporary', turbines that would not be used for over 365 days, which some counties argue do not require any sort of permitting.



## April-July 2025 — Aerial imagery reveals a total of 35 xAI gas turbines at Colossus 1

Suspicious over rapid construction, the SELC charts a volunteer flight and takes thermal photos, showing 33 gas turbines in operation, capable of producing 421 MW. This is more than double xAI's permit application.

In a nearby County Commissioner's meeting, residents share their stories of chronic respiratory issues.

Based on the imagery, SELC argues that xAI did not account for all turbines to keep the site below major-source permitting, in violation of the Clean Air Act.

In July, the Shelby County Health Department exempts the unpermitted turbines as 'nonroad engines' — which don't require permitting. Critics argue this shielded xAI from the consequences for roughly a year of unpermitted operation.

State agencies like MDEQ do not seem concerned with enforcing environmental standards, in fact they appear to provide their support so data centers can work around regulations— without scrutiny. This means local legal and activist groups must do their own investigations.



## March–May 2026 — MDEQ approves 41 permanent turbines for Colossus 2

In March 2026, despite strong public opposition, MDEQ granted xAI an air permit for 27 temporary gas turbines. These turbines would sit across the border in Southaven, Mississippi but power Colossus 2 in Tennessee.

Critics noted the approval relied on xAI's air-quality models.

**SELC and Earthjustice sued xAI for violating the Clean Air Act, arguing that regulators were providing unusual leniency:**

- Shelby County classifying as a minor source
- Shelby County treating turbines as exempt “nonroad engines”
- MDEQ in Mississippi accepted the turbines as mobile because they sat on a trailer
- Neither investigated nor enforced the number of permitted turbines

## Conclusion

Despite local failures, we are now aware of the most important parts of a data center's legal life cycle, from idea to a fully permitted operational data center.

Its clear hyperscale developers have been abusing legal loopholes in a major way and current regulation and enforcement is not adequate. Yet this issue brings the whole concept of local government and environmental protection agencies into question, given their lack of care with projects so obviously rejected and unwanted by residents.





12

# How To Spot and Stop a Data Center



Based on everything we've learned in the legal life cycle section above, we know there are a few crucial points in a data center's life cycle where residents can push back.

**Concerned individuals need to raise community awareness before a data center project receives approvals.**

### - Preparing for the Search

If searching for a known data center in your area, it should be easy to track down the address by reading local articles or by searching online. You can also find the address of projects on **Data Center Map**, do note this is an industry resource with no affiliation to our AI Data Center Map focused on awareness and ecology.

## 1. Site Acquisition


Your search starts with the property records, which can be found in two local offices:

- **County Appraisal District/County Assessor:** This office determines the value of all properties in the county and has records of their characteristics, and development over time. It can be useful to track hyperscale infrastructure development through yearly appraisals
- **County Clerk/County Recorder:** Preserves personal and property records including deeds, mortgages, liens, easements. The official record of who owns what. The deed records the transfer of ownership from the previous owner to a hyperscale developer.



### TIP

Search online for the name of your local county plus the office you're looking for to find their web search. (e.g "Hays County Appraisal District")



The County Appraisal District site for Hays County

Some county search tools are built better than others, at minimum you can search using the following:

1. **Exact address of the property** - this may be found on local news articles or through a web search
2. **Name of the property developer** (less reliable as shell corporations are typically used)

## Tweaking Your Search

Each project is different and you'll need to adjust your search. In our example, the exact address 904 Francis Harris did not show up in the search, but we were able to find it using the name of the owner.

In other cases, its very common for developers to purchase the land through an unassuming shell company, and if you can't find the name you're better off using the property's address.

The most important thing is to be resourceful and persistent.

Page 1 of 1 - Total: 1 (0.038 seconds)

Customize View Year: 2026 Page Size: 25 Download Results

| 1            | 1 - 1 of 1 items |      |                       |          |                                             |             |
|--------------|------------------|------|-----------------------|----------|---------------------------------------------|-------------|
| Quick Ref ID | Nbrhd. Code      | Type | Owner Name            | Owner ID | Site Address                                | Appraised   |
| R70024       | C-SM-35OBR       | Real | HIGHLANDER SM ONE LLC | O0134766 | 900 FRANCIS HARRIS LN, SAN MARCOS, TX 78666 | \$6,985,920 |
| 1            | 1 - 1 of 1 items |      |                       |          |                                             |             |

## - Appraisal District Findings

The appraisal district is a great first stop for information about a hyperscale project:

- **Property Details:** The type of property, its exact location and who owns it
- **Owner Information:** Who or what company owns the property, and their mailing address
- **Yearly Appraisals:** The market value of the property year over year, and current infrastructure on the property including buildings and land



## - County Clerk Findings

The county clerk maintains important property records including Deeds, which formally record the transfer of ownership, and may also contain financing information and other property details

The transfer dates indicate when the project was initially launched, far before public announcement.

| Document Type |  |  |
|---------------|--|--|
| DEED OF TRUST |  |  |

| Recording Information               |                                                  |
|-------------------------------------|--------------------------------------------------|
| <b>Document Number:</b><br>25023997 | <b>Recording Date:</b><br>07/01/2025 08:49:45 AM |
| <b>Number Pages:</b><br>26          | <b>Recording Fees:</b><br>\$121.00               |

| Book Page |        |      |
|-----------|--------|------|
| Book      | Volume | Page |

| Names                                    |  |
|------------------------------------------|--|
| <b>Grantor:</b><br>HIGHLANDER SM ONE LLC |  |
| <b>Grantee:</b><br>R BANK                |  |

**Deeds of Trusts**, is a property acquired with the help of financing, and can reveal the borrower's parent entity. Loan amounts can reveal project cost.

**Surveys / Plats** show easements or special features which can hint at power infrastructure.

### - Parcel IDs

You may see a Parcel ID or a reference ID in your search, which makes it easier to cross reference the project's documents on other parts of the legal lifecycle (permits, environmental, zoning).

#### TIP

Cross referencing the shell company names across county clerk documents can potentially lead to the lender, or even parent guarantor — revealing the true property developer.

## 2. Tax Breaks

Any special tax abatements or deals at the city or council level are discussed in court agendas. By keeping up to date with new meeting agendas and minutes, you'll be able to see if there's a data center in the works.

Look for any interesting items of business related to third party property development, tax abatement, or permitting you don't recognize. There should be a link to any supporting documentation.



## Comptroller

The state comptroller's registry may contain tax subsidy or abatement reports. The Texas comptroller has their state-specific Chapter 312 Abatement Search Tool, which lets you search by company or property name.

Remember, only 20 states publicly report their tax subsidies — so you'll need to determine if your state has a searchable registry.

If not, you may try the **Good Jobs First Databases**, which bring together 3 million records from federal, state, and local sources. This data is combined into subsidy and tax break trackers as well as other useful transparency tools.

### 3. Zoning and Land Use

Within the county planning and zoning commission, city council, or county commissioner's public hearing minutes you may discover industrial zoning, special use permits, and city plan amendments for data centers.


### 4. Power and Interconnection

To get an overview off the U.S power grid, visit **Interconnection.FYI**. Interconnection FYI tracks power generation requests to major U.S grids.

#### U.S. county map of active generation interconnection requests, 1995-2026

No. requests    Total capacity  
**8,188**    **1.76 TW**  
For active generation requests

DATA LAST UPDATED JUN 1, 2026, 10:36 PM CDT



Most grids publish records around power generation requests, but not demand requests — these are discussed in special committees.

**This requires one to dig through the committee and working groups hearing notes:**

1. Determine which power grid serves your area of interest (CAISO, ERCOT, PJM, NYSIO, MISO, etc...)
2. Find the committee or working group that is responsible for interconnection requests and/or large loads interconnections, go to their
3. Attend the meeting in person, call in, or read through the published minutes to see the orders of business

## **Texas's ERCOT Large Load Working Group**

You can visit ERCOT's large load working group here (<https://www.ercot.com/committees/tac/llwg>), reading through their previous meeting minutes we see they discuss the large interconnection queue, and power studies.



## 5. Environmental

To get ahead of hyperscale driven air pollution, you can check air and water permits on your state or county level environmental agency.

Depending on the state you may find different resources. A good starting point is **The Environmental Database**, which has quick links for every state environmental agency.

### Searching Nationally

The EPA has systems like ECHO (<https://echo.epa.gov/>) which track compliance to federal regulations like the Clear Air Act and Clean Water Act. These systems are supposed to contain history, inspections, and violation — but are typically delegated to states and even counties.

Although in some cases, hyperscales do not appear on the ECHO map — the map view is still worth using as you can see permits in any given area. Cross referencing the data centers location on the ECHO map will give you useful hints to the permit representatives.

For Meta's Hyperion, the air permit was registered by the partner power provider for Meta, Entergy.

## Searching By State

State environmental agencies are delegated to uphold the federal Clean Air Act, often you must search state level environmental agencies, or even county level for environmental records.

1. Start this process with the state agency, if nothing is found, repeat the process by searching for the local county's air permitting page
2. Once you've found the agency, search through their air or environmental permit database.
  - a. Search by **street address** or **Parcel ID**
  - b. If it allows for a map search, scan over the location of interest
3. Hyperscale data centers will be operating either emergency gas turbines or auxiliary generators, these will almost always be permitted as a **minor source**
4. If no state or county level reports are found, look for the public-notice and comment archives

# Let's use this process to pull up air permits for Meta's Hyperion in Louisiana.

1. Find the state agency  
**(Louisiana Department of Environmental Quality)**
2. Find the **air permit status search** or similar database to find the facility's ID
  - a. Filter by the data center's county  
(Hyperion was located in Richland)
  - b. Cross reference with any business names or energy partners associated with the project (Entergy Louisiana built a special power plant for Hyperion, Franklin Farms Power Station)
  - c. Keep note of the facility ID
3. Find the state environmental agency database, central registry, query or program search, or document management system. (Louisiana uses an **electronic document management system**)
4. Use the facility ID you found in step 2 to search for permits

## 5. Find any documents labeled permit to find direct information plans and specifications for hyperscale power generation

Basis for Decision  
Entergy Louisiana, LLC – Franklin Farms Power Station  
AI No. 245286  
Permit Nos. 2460-00034-V0, PSD-LA-862, and 2460-00034-IV0

### C. Description of Facility

The Franklin Farms Power Station will be a 1500 megawatt (MW) power generation station consisting of two combined cycle generating units. Each unit will be comprised of a Mitsubishi Model M501JAC gas turbine, a supplementary-fired (duct burner) heat recovery steam generator (HRSG), and a steam turbine. Other emission sources will include standby generators, diesel firewater pumps, natural gas dewpoint heaters, condensate tanks, wastewater, fugitive emissions, General Condition XVII Activities, and Insignificant Activities.

Although searching for air permits isn't a straightforward process, seeing these plans ahead of time can show specifics like the total power generation, models of generators, and even emissions outputs. Local activists can use this to craft a strong case against the next local hyperscale data center.

## 6. Public Records Requests

If you can't find what you're looking for in the county, you may attempt a state level Public Records request. Every state has an open records statute that requires custodians to respond within ten days of a specific request. Using what you've learned here, you can request specific documents, permits, or filings for data center projects. Most jurisdictions must respond in ten days.

### State Records Requests

The **Open Government Guide**, contains a list of every state's open records and meeting laws. For each state, they share who can request records, which organizations and records are subject to public, fees, and any exempted information.

### Federal Documents

The Freedom of Information Act (FOIA) gives any person the right to request records from the federal government. Although details on data centers are typically state level, knowing how to submit FOIA requests is always helpful. The first place to start would be the **Freedom of Information Act Wiki's FOIA Basics Guide**.

The **FOIA Wiki** is another great resource to give you background on FOIA, its history, the process of making a request.



# 13 Legal Climate



## Emergency Grid Regulation

In 2025 three major regional grids proposed emergency regulations to accommodate the massive influx of data centers, which account for the majority of large load interconnection requests.

These new regulations provide data center developers more leeway to connect to the grid, including expedited interconnection programs, where interconnection studies are fast-tracked and hyperscale developers are invited to “BYOG” (“Bring Your Own Generation”).


## States Giving Billions In Tax Subsidies

State level sales tax and property tax exemptions significantly cut down costs for construction and materials, servers, cabling, data storage, security, power and cooling infrastructure, and operations centers. The largest scale hyperscale projects can save hundreds of millions — if not billions.

According to non-profit watchdog group **Good Jobs First**, only 11 states disclose the recipients of data-center tax exemptions. The names listed are usually shell companies, and states do not demand public transparency on the companies that ultimately benefit.

Only 5 states (Illinois, Indiana, Nevada, Ohio, and Washington) disclose the dollar amounts received by specific companies.

And no states report on the amount of jobs promised versus the jobs created — there's no real way to measure whether hyperscale developers are living up to their promises.








**The amounts of these subsidies can be astounding. Texas lost over \$1 billion to data center subsidies in FY 2025—the most expensive subsidy for any industry in any state.**

**In addition to state tax subsidies, data centers may receive local and regional tax subsidies. The combination of these subsidies boosted Virginia’s aggregate subsidy from \$732 million to \$928 million in FY 2024.**

# States Are Handing Data Centers Hundreds of Millions in Tax Breaks

Sales and use tax exemptions allow data center projects to avoid paying taxes on major development and operating costs, from servers and cabling to power infrastructure and cooling systems.

## Largest Reported Data Center Tax Breaks



Note: Most recently available data; years vary by state. Figures reflect reported subsidy costs and may not include all local or regional incentives.



### Texas passed \$1 billion in one year

The state reported more than \$1 billion in data center subsidy costs in FY 2025.



### Virginia's subsidy bill was even higher

A separate study found \$928.6M in combined tax exemptions in FY 2023.



### Illinois surged by 3,600%

Reported revenue loss grew 3,600% over four years, reaching \$370.6M in 2023.

[aidatacentermap.org](http://aidatacentermap.org)



**Because there are no controls or caps placed on the subsidy programs, costs have exploded. In Illinois, the reported cost in lost taxes from the subsidy program grew 3,600% in four years. In 2023, Illinois was out a cool \$370 million in terms of total taxes lost.**

Meta is building one of the largest data centers in the world in Louisiana, yet it will not have to reveal how much in state taxes it is avoiding. Only four states (Illinois, Nevada, Missouri, Washington) show tax breaks given to each of their data center projects.

Given that data centers receive such unfair advantages through these tax subsidy programs, it is no wonder they are rushing to get projects approved and operational. Meanwhile, the explosion in grid interconnection and infrastructure are driving up prices for those living and working nearby.

This perfect storm has states & counties second-guessing their tax programs.

## Data Center Moratorium Bills Are Going Dead

Thirteen states now have proposed moratorium bills, which involve pausing approvals for construction and incentives for data centers.

But this effort fell far too short. None have passed. The Maine legislature approved an 18-month pause, but that bill was vetoed by Governor Janet Mills. Other states, too, have been experiencing fierce opposition from power plants and developers.

Within local governments and counties, the story sounds a little better.

The **US Data Center Moratorium Tracker** is counting 69 local moratoriums across counties, cities, and towns. Of those 69 moratoriums, 50 are active.

Still, these local moratoriums often fall short. They grandfather existing data center developments. And they only propose one-year pauses, even though data centers are being planned through 2038. Oklahoma has proposed a moratorium that would apply until 2029, although it is unclear if it will be passed.

These moratoriums may simply be lip service—a way to show that politicians care about their constituents—without making any meaningful change.




14

# Gigascale Projects Transforming the US







Let's take a tour of some of the largest and most controversial hyperscale projects in the United States. Hyperscalers use 100 megawatts (MW) of power minimum (enough to power approximately 100,000 homes).

But the term “hyperscalers” falls short when it comes to describing the **1 GW+** data center projects now seen around the United States. We term these huge data centers “**giga-scalers**” or **gigascale**, which use power equivalent to a million homes per GW.



Developers of gigascalers use underhanded tactics to establish themselves. They minimize their public presence until their project is approved by the city or county by using shell corporations — before the public has a chance to organize.

In early stages, they prospect multiple towns in poorer areas, using their negotiating leverage to receive the best tax breaks, abatements, and commitments to power and municipal utilities, while minimizing regulatory risk and costs.

In this section, we observe their *modus operandi* and conduct a post-mortem of the community—after the economic shock of surging rent prices, homelessness, and agricultural damage.

There are two main groups driving these gigascale projects: big tech-managed sites, which vertically integrate their infrastructure, and energy developer-led projects, which are developing sites with power and water, hoping that tech firms will rent their facilities. Often energy developers have no planned occupants. They create massive shock waves in the community based on pure speculation.

To shine some light on the dark approaches used by data center developers, let's look at six of the biggest gigascale data centers in the United States and their path from planning to being put in operation.

#6

## “Ran Too Fast”: Project Rainier

New Carlisle, Indiana

Less than one year after its announcement, Amazon’s Project Rainier is operational as of October 2025. The \$11B project is the largest in Indiana’s history, and the largest power draw in the entire state at 2.2 gigawatts of electricity.

The site is an important part of Amazon's AI strategy; 500,000 "Trainium2" chips, designed by Amazon run at this facility. Its reportedly the largest cluster off non-NVIDIA chips in the entire world. They're used to operate the large language models of Anthropic, which Amazon has invested \$8B in. These chips aren't as powerful as their NVIDIA counterparts, but they are designed specifically to be used in larger numbers for lower-cost AI inference.

Given the greenlight by Indiana, Amazon received sales and use tax break for all projects in the state that could save the company billions. According to the Indiana Economic Development Corporation (IEDC) Amazon reported \$560M in sales and use tax exemptions in 2025 alone.

Despite unusual support from the local government, the project comes with controversy as the facility sits on top of marshland and wetlands.

To make solid ground, Amazon applied to fill nine acres of these native marshes with construction material and route 5,000 feet of streams in the area.



The project was found to be in violation when one of the company's contractors began dumping construction material without approval, causing a notice to halt work by the Indiana Department of Environmental Management.

## **At a town hall meeting, local farmers called out the recklessness:**

**“The bigger picture is, they are taking water out of that aquifer and not replacing it. How it gets replaced is the wetlands. So instead of allowing them to destroy what's currently there, we should be encouraging them to expand it [aquifer wetland recharge zone] tenfold to replace all this water,”**

**Said Gredoms[,], who lives and farms in new Carlisle.**

To remediate the situation, Amazon sought to purchase wetland and stream credits; which is encouraged by the Clear Water Act as a way for companies to offset their environmental impact.

After resubmitting their application, construction resumed quickly. Amazon carries on with its construction working towards their goal of 30 buildings across 1,200 acres by 2030.

The NYT reports the facility will use millions of gallons of water every for cooling, and requested permits for 'dewatering' the site. This is a process used to remove water from an area to lay cable and build underground infrastructure.

The permit requested to pump 2.2 million gallons an hour for 730 days. If we assume full use of the permit, that would be 38.5 billion gallons removed from the wetlands in total. Its unclear if this was the total amount pumped, but officials are investigating if this process caused dry wells in the area.

As we've learned, water is the lifeblood of local ecosystems. To have groundwater pumped out from wetlands at this scale is a major shock, one that may be felt for years. In our view, you can't undo that kind of damage, no matter how many wetland credits you buy — the regulations meant to protect the environment are falling far too short.



#5

## “Not Fair”: Fairwater 1

Mount Pleasant, Wisconsin

Touted as the “world’s most powerful” AI data center, Microsoft’s Fairwater facility in Mount Pleasant, Wisconsin, is outfitted with hundreds of thousands of NVIDIA cards operating as a single cluster. The site, which sits on 315 acres of land, is 1.2 million square feet across three buildings.

Microsoft was proud to announce the facility’s closed-loop water cooling, which they claim will have zero water waste and evaporation.

The coolant is supplied once and circulates through each server, absorbing heat from the AI chips.

The coolant then travels to one of the largest water-chilling systems on the planet, which dissipates the heat with 172 20-foot fans.


Microsoft claims that in one-fill, the closed-loop liquid-cooling system uses what 20 homes use in a year, estimated at roughly 2 million gallons (assuming a home uses 300 gal/day). This amount would be replaced once every 6 years.



Rack level liquid cooling



Despite this feat of engineering, we're not convinced of its minimal ecological impact. They claim the closed-loop cooling only accounts for 90% of the facility's compute, with the final 10% on air cooling that switches to water cooling during the hottest days.



Aerial view of the liquid cooling system

This means that the facility can pull water directly as needed, at any point in time. As we read in the sustainability section, facilities do not even meter their water usage. Its unlikely we will ever true measurements of water or power the site uses, although local grid operator filings show that the facility is rated to pull 1,500 MW.

Some 20 miles north of Fairwater Wisconsin in Mount Pleasant, Microsoft hatched a more secretive project, Project Nova.

Project Nova would sit on a 244-acre site in Caledonia, abutting Lake Michigan. Unlike Mount Pleasant which was announced publicly and marketed, Project Nova was hidden from the public as Microsoft attempted to rezone the land in September 2025.

In the rezoning hearing 49 residents spoke. Of those, 39 opposed the plan. They cited three main concerns: environmental harm, the possibility of taxpayers having to foot the electricity bill, and worry that their animals would suffer from the noise.

Despite the opposition, the commissioners voted 5-to-2 to approve the rezoning of 244 acres from to light industrial/manufacturing. This would move onto a final city council vote, but before that vote could happen, Microsoft withdrew from the project.

How did this push back start? It all started with observing a simple yard sign that said “rezoning application”. After realizing what was at stake, 2,000 Caledonia residents signed a petition against the data center. If they had missed it, the project may have been approved without their knowledge—or input.

Residents must be vigilant stewards for their communities.



#4

## “Gate To Ruin”: Project Stargate

Abilene, Texas

The \$500-billion Project Stargate is a private joint venture announced by President Donald Trump at the White House Press Conference on January 21, 2025—the second day of his second administration.

Standing alongside Trump were some of the project’s collaborators: OpenAI CEO Sam Altman, Oracle Chairman Larry Ellison, and SoftBank CEO Masayoshi Son. (The fourth partner is UAE-based investment firm MGX.)



The plan is for Project Stargate to invest up to \$500 billion in AI infrastructure in the US by 2029. Phase 1 saw seven facilities across 5 states with a total power capacity of 7 gigawatts.

The Abilene, Texas, campus is the venture's flagship site, spanning up to It spans more than 1,000 acres of land and will house ten data centers.

Abilene is favored for many reasons: lots of undeveloped space, alternative power, and a 1,200 acre site that did not require rezoning, meaning that locals couldn't meaningfully stop the project.

The local government has been driving the narrative of local acceptance, Abilene Mayor Weldon Hurt sings the praises of the project:

**“It will impact the rest of the economy—our restaurants, our homebuilders—with that many new people coming in and taking these jobs.”**

Mayor Weldon Hurt  
**[Texas Observer]**

To fast-track this project, OpenAI and Oracle contracted directly with the Abilene City Council without first informing the public. Residents didn't hear about the project until it kicked off.

Mayor Hurtand and city officials put together a 85% property tax abatement, which requires Crusoe, the company responsible for operation to maintain 357 permanent jobs — which aren't required to go to locals.

An analysis from the Development Corporation of Abilene (DCOA) estimated a positive \$4B economic impact over 20 years, but its unclear how it would reach those heights with only 357 permanent jobs; and DCOA has refused to share their methodology.

When City Council members held an unorthodox special meeting to vote on the tax abatement, only two members of the public showed up. Naturally, the data center was approved.

## Is this small-town politics at work?

The answer may lie in what has happened since the City Council vote. Stargate has opened the gate to all kinds of chaos in Abilene. Six thousand construction workers, along with all the heavy machinery required to build big data centers, flooded the town, endangering roadways and creating unease.

**Farmlands and pastures are now transformed into construction zones, its dangerous for animals to graze with construction traffic.**

**New workers have caused rents to jump \$1,000 in a single year, according to Zillow data; locals report rents double and tripling.**

**Habitat for Humanity began to see entire families showing up at homeless encampments, unable to afford the rising rents.**

**As for the neighborhoods within half-a-mile of the site, their lots and homes have plummeted in value.**

**“The price of progress is  
always hard”**

**Says Mayor Weldon Hurt.**

**“If there’s a construction  
project, it’s inconvenient. And  
this one’s just huge, which  
makes it hugely inconvenient,  
especially if you live out  
there.”**

**Hugely inconvenient,  
or a humanitarian crisis?**



#3

## “Colossal Problems”: Colossus from xAI

Memphis, Tennessee

No doubt about it, xAI’s Colossus facility is a burden for the residents of Memphis, Tennessee.

When the project opened in 2024 with one site, “Colossus 1,” was announced as one of the largest data center in the United States breaking records for size and speed, operational within 18 months.



To power the Colossus 1 site in Memphis, xAI brought in a dozen unpermitted gas turbines in June 2024, growing to 35 turbines by Spring 2025. xAI had only filed permits to use 15 turbines for their power generation.

When the Southern Environmental Law Center (SELC) decided to fly over the facility in April 2025, it discovered that 33 of these 35 turbines were up and running.


Intense community pushback had no effect. SELC said it intended to sue xAI. After that, XAI removed some of the gas turbines but reportedly continued to operate additional unpermitted turbines on the site.

In a glaring failure of government, the Shelby County Health Department gave xAI a pass for the unpermitted turbines by providing a non-road-engine exemption.

After this controversy, xAI officials said they planned on “copying and pasting” their power strategy. The Tennessee project quickly expanded, with xAI acquiring two additional sites: (1) a second Colossus, “Colossus 2,” in nearby Southaven, Mississippi, meant to house an additional 550,000 NVIDIA GPUs (and intended to go live in April 2026), and (2) land to house natural gas turbines to power the facilities.

Now, in May 2026, the SELC and other organizations are appealing the permit and requesting a preliminary injunction to halt the operation.


XAI claimed its emissions were low or limited, but thermal imaging appears to show many turbines actively running, suggesting **the pollution source was larger than the public had been led to believe.**



According to Laura Thoms, the  
Director of Enforcement at Earthjustice:

**“xAI is running a de facto power plant without an air permit, without necessary pollution controls, and without regard for families living as close as a half mile away. We won’t let the company disregard the law and the health of Mississippi residents.”**

Meanwhile, Memphis (TN) and Southaven (MS) residents suffer from some of the worst air quality levels in the country. Not surprisingly, they complain of chronic respiratory issues as that air quality continues to worsen.



**After Colossus 1 was up and running, at the request of TIME Magazine researchers at University of Tennessee, Knoxville performed an air quality analysis and found a 3% average increase in nitrogen dioxide (NO<sub>2</sub>) concentration and increases of peak concentrations by 79% in the vicinity and 9% in the neighboring town of Boxtown.**

Elevated NO<sub>2</sub> levels harm the lungs and aggravate respiratory diseases, causing asthma attacks, coughing, and difficulty breathing. NO<sub>2</sub> can also lead to smog and acid rain.

Meanwhile, Colossus 1 is contracted to use 1.3 million gallons per day from the municipal water supply—an amount greenlighted by the president of Memphis Light, Gas and Water.

After community pressure was exerted by organizations like Protect Our Aquifer (POA), xAI committed to using recycled wastewater, although no price has been finalized. The water would only take care of only 20% of Colossus' needs. The rest would come from municipal water, which is fed by the Memphis Sand Aquifer.

Memphis is one of the largest cities primarily dependent on groundwater. Now that there are major demands from heavy industry (including the data center) being made on that groundwater, pollution on the surface could jeopardize Memphis residents' drinking water.

**The situation in Memphis reminds us to never put our trust in companies on their word and to scrutinize and verify things for ourselves.**



#2

## “Hyper-risky”: Meta’s Hyperion

Richland, Louisiana

Hyperion is Meta’s largest data center to date. It spans four-million-square-feet in Richland Parish on former farmland. Announced in 2024 as a \$10B project, its now exploded to \$200B in project costs, with a power capacity of 10 GW when it goes into service in 2030. It will be tasked with training Meta’s Llama AI models.



**Meta was able to obtain excellent terms from the local government, saving the company billions in costs. For example:**

- Meta is **exempt from any state and local sales and use taxes** on data center equipment, saving it a total of **\$3.3 billion** in taxes on qualifying equipment purchases.
- Meta received fast-track approval for **ten new permanent gas power plants** from Louisiana energy provider Entergy.
- Entergy's other customers have to pay for a 60-mile transmission line to the Meta data center at a cost of \$470 million.
- If it invests at least \$10 billion in Hyperion, Meta is required to pay a mere 20% of its property taxes every year.

**Critics argue Meta pulled a financial bait-and-switch on the Louisiana Public Service Commission (PSC).**


The Hyperion deal was approved only because Meta promised that ratepayers would not have to pay for the new gas plants. The same day the vote passed, Meta restructured the project. It started a joint venture with private equity firm Blue Owl Capital, shifting 80% of its ownership into a new company, Beignet. This new company borrowed **\$27 billion** to fund the data center, which lets Meta walk away from the data center after four years without any penalties.

Because the gas plants have a 30 year life span, consumer advocates like Earthjustice, the Alliance for Affordable Energy, and the Union of Concerned Scientists warn that if Meta exits early, then the costs of the gas plants will be footed by ratepayers.

Louisiana Public Service Commission which regulates utilities refuses to investigate the restructuring, leaving residents in the region vulnerable.

Hyperion is registered to consume more than **23 million gallons of water per day**, according to state records. This would be roughly 8.4B gallons every year.

However, Meta says not to worry: Hyperion will actually only use **1.5 million gallons a day** (or 600 million gallons a year) due to its closed loop system which they claim will require no water during the cooler months.



The water is drawn from the Mississippi River Alluvial Aquifer, which supplies farmlands in six states, which together have a **\$12 billion agricultural economy**. The water supply from the aquifer is relatively abundant because it is recharged by the rains in the river valley.

Despite Meta's promises—including its vow to voluntarily submit annual water usage figures for the first five years—there is no regulatory body monitoring Meta's water usage or the aquifer conditions over time.

**“There’s no actual scrutiny on the use of public resources.”**

Tyler Gray,  
Secretary of Louisiana’s Department of Conservation and  
Energy (January 2024 to September 2025)

The nearby town of Delhi will supply water for two of Meta's electricity plants, and residents are concerned about their own water running dry.

In the past, Louisiana officials have sacrificed the state's environment for big business. The clearest example is the 85-mile stretch of land along the Mississippi River between New Orleans and Baton Rouge, home to some 200 fuel and petrochemical plants built over decades.

This area has been known since the 1980s as “Cancer Alley” due to the elevated risk for cancer, verified by researchers at top universities including John Hopkins.

Hyperion follows these same patterns; an industry with massive pollution given generous tax advantages. It is prudent that we organize locally to prevent a repetition of Louisiana’s deadly mistake.

**“One of the first things that I always say when I go to communities is, ‘Our opponents always think that we’re not smart enough or passionate enough or rich enough to stop them. But in my career, I have seen how everyday people can prove them wrong.’”**

Tim Cywinski, Communications Director  
for the Sierra Club Virginia Chapter



#1

# “Project Bull” Fermi America’s Project Matador

Texas

Deep in the Texas Panhandle sits Project Matador, one the most ambitious energy and data projects in the entire world, with a total build-out estimated at hundreds of billions.

Project Matador’s is being built by Fermi America, a brand-new company founded by former Texas Governor and U.S Energy Secretary Rick Perry.

**Fermi America has never built a single data center before. Yet the company is so audacious that if they fulfill their vision, they may build the largest data center in the entire world.**

**The 7500-plus-acre site has set a target to produce up to 17 GW of power—as much as the entire state of New York uses on a single day or the equivalent of 12-to-14 million homes.**

Project Matador sits atop one of the largest natural gas fields in the United States. Two clean air permits suggest that Fermi will use natural gas to power 11 GW of total power. The rest of Project Matador's energy needs will be supplied by four Westinghouse AP-1000 nuclear reactors, which are scheduled for completion by April 2031, and a large solar grid.

**“Ultimately, when completed, it’ll be about 50% natural gas, 40 or 45% nuclear, and the remainder solar.”**

**Griffin Perry,**

**son of former Texas Governor Rick Perry and co-owner, with his father, of Fermi America**

These are long-term plans, however as of June 2026 Fermi only 6 GW has preliminary approval, and only 2 GW of generation has been secured so far.

Both the natural gas and nuclear plants will need tremendous amounts of water for cooling. The agreement Fermi reached with the City of Amarillo allowed for up to 2.5 million gallons a day over a twenty-year contract. However, Fermi has since said it may potentially need up to 5.5 million gallons a day and expects to sign a new agreement with the city.

Thankfully, Fermi’s water usage will be measured by a meter that the city can read. But, if water exceeds the agreed-upon amount, there are no financial consequences built into the contract.

Last October, in a presentation to the Amarillo City Council, Fermi said if denied the necessary water, it would purchase eleven nearby cornfields and pump directly from the Ogallala Aquifer.

Ogallala is a major aquifer underlying the High Plains region supplies parts of eight-states with water. and has long-since suffered from drought. The farmland it contributes \$35B.

As we read in the water rights section, Texas has 'absolute dominion' or 'rule of capture' water rights, as long as you can pump it — you own the water.

Texas does have Groundwater Conservation Districts in play which create and enforce water allocations. However we've seen in past cases of violations and fines, we've seen defendants turn around and sue the Groundwater Conservation District, while continuing to pump.

Many times the fines are excused, or fined at a few dollars per thousand gallons.

If Fermi abuses their usage, its subject to the Panhandle Groundwater Conservation District which covers the site. However, would this local GCD be able to take on a billion dollar gigascale developer? If not, we could potentially ruin the regions agriculture.



But that's just the beginning of the controversy around Fermi. A year into the project, the site still hasn't decided what to do with nuclear waste. Until Fermi finds a storage partner, it plans to store the cooling nuclear waste onsite.

Moreover, before the nuclear plants become operational, Fermi intends to bring in 93 portable gas-fired turbines. That plan received 279 comments in a public hearing—mostly locals objecting to Fermi's air quality permit.

Despite those objections, Fermi still received approvals for a total of 6 GW in natural gas air quality permits from the Texas Commission on Environmental Quality (TCEQ). That's the second-largest air permit in the United States. Fermi filed a subsequent permit, bringing the total to 11 GW of natural gas power.

Thankfully, this project seems like its dead in the water. At the end of 2025, Fermi's first tenant for the data center pulled out. A class action lawsuit followed, accusing Fermi of overstating tenant demand.

The project is holding off on further construction until a tenant is found, and according to Cleanview's satellite imagery, this looks to be the case. As of April 2026, when Fermi had originally said it would complete construction of one-million square feet of facility space, the construction site looks the same.

Project Matador got hit by a bull and we're grateful for it.



## Section Review

We've learned a lot from real-world gigascale case studies:

### How They Get Approved

- Gigascale developers discuss plans in private with city and county officials.
- If a public hearing is required, notices are minimized, posted on in-person signs and placards instead of notifying all residents
- Special meetings are held to discuss projects and subsidies
- Unless residents are proactive, or a local conservation group steps in, giga-scale projects are quickly approved by public officials before anyone can react
- Data center developers make immense profits from tax breaks and sweetheart deals.

# Impact of Gigascale Projects Construction & Operation

- Arrival of thousands of construction workers, heavy equipment, and construction materials
- Traffic delays are caused by construction traffic, road blocks, and hazardous conditions.
- Surging rent prices for temporary housing of construction workers may cause homelessness.
- Huge drops in neighboring property values due to the noise and pollution
- Massive power draws can cause energy bills to double, triple, or even quadruple.
- Air pollution from onsite gas turbines and power generation causes health and environmental problems
- Heat released into the atmosphere is felt by everyone
- Nuisance noise from daily operations includes a constant loud humming.
- Nearby wells may go dry
- There is a risk of the project financing failing, leaving the other energy customers to foot the bill.

15

# Dangers To Living Things



## Air Pollution

In February of this year, researchers at Virginia Commonwealth University (VCU) studied the air in Northern Virginia, a region known as “Data Center Alley” for its high density of hyperscale data centers.

They looked at both maximum permitted air pollution and *actual* air pollution emitted by 138 Northern Virginia data centers between 2015 and 2023. In that eight-year period, they found major spikes in harmful pollutants:

- Carbon monoxide emissions jumped 196%
- Nitrous oxide emissions rose by 111%
- Dangerous particulate matter jumped by 139%

These regional data centers across Arlington, Fairfax, Loudoun, and Prince William counties account for 4% to 7% of the total tons per year of emissions that they are permitted to release — primarily from their 4,000 diesel backup generators.

The emissions from backup generators minimally operated already exceed the emissions from nearby power plants. The backup generators are ran for maintenance, testing, or short-term emergency outages, however if there was an extended outage, electric grid operators could request that data centers run on fully their backup power.

Running at full capacity, these generators would produce ten to twenty times more pollution—one third of total nitrous oxide emissions and half of carbon monoxide and particulate matter pollution.

In a state that already deals with coal plants and heavily congested highways, this is no joke.



Damian Pitt,  
PhD and one of the VCU study's three authors.

**“You’d rather be living next to a natural gas power plant than living next to a whole bunch of data centers.”**



## Health Impacts

Compared to diesel generators, natural gas generators produce much cleaner exhaust. Still, natural gas pollutants include carbon dioxide (CO<sub>2</sub>) and nitrogen oxides (NO<sub>x</sub>).

Diesel generators, by contrast, emit more harmful pollutants and at higher levels, from carbon monoxide (CO), nitrous oxides (NO<sub>x</sub>), and sulfur dioxide (SO<sub>2</sub>) to particulate matter (PM 2.5) that is 2.5 micrometers or smaller.

**These outputs are linked to respiratory and cardiovascular diseases, cancers, neurological disorders, and conditions such as stroke, dementia, diabetes, and asthma. They can also result in reduced fertility, fewer live births, and increased miscarriages.**

Particulate matter at 2.5 micrometers is small enough to penetrate the natural filters in our airways and move deep into the lungs and bloodstream and even pass the blood brain barrier, leading to cognitive decline and neurodegeneration. In addition, the VCU study found

ozone precursors and toxic air contaminants (black carbon, benzene, and formaldehyde, for example) emitted from the diesel backup generators.

Many new data center projects in grid-constrained area are front-loading power sources with diesel and natural gas generators. These are used as transitory & emergency power while permanent power sources are built.

Electric grids are forecasting extremely high demand through 2030, some grids like PJM find that they can't provide enough reserve power in the coming years to meet demand.

This may lead to instability and grid emergencies being declared, which in turn will require the use of gas and diesel generators to power data centers.

However, in an emergency, the generators may run continuously to keep the lights on; leading to unprecedented pollution levels.

## Data Heat Island Effect

A coalition of academic researchers in the UK, France, Italy, Singapore, and China has discovered what they call a “data heat island effect”.

Using 6,733 data points for land surface temperatures, the researchers were able to measure the heat increase around hyperscale data centers.


These facilities consume massive amounts of electricity to perform calculations. A single professional AI card can perform a trillion-trillion calculations per second on its 2-billion transistors. Each chip generates so much heat, it can be likened to the surface of the sun.

All this heat across thousands of AI chips must be dissipated. Liquid cooling or air cooling take the heat and release it into the air, creating the aforementioned “data heat island effect.”

The researchers determined that this effect can first be felt within a month of a data center’s operational start. They measured an average increase of 2 degrees Celsius (3.4 degrees) within a kilometer to hyperscale data centers.

This increase in temperature extends up to 10 kilometers outside the hyperscale facilities but drops in intensity

fairly linearly, with a 30% intensity at 7 km out. In extreme cases, temperature increases of up to 9 degrees Celsius were measured.



The data centers studied were outside of densely populated areas, to better isolate the data heat island effect from the well-known urban heat island effect, caused by high-density population infrastructure and a lack of heat-absorbing soil and plants.

The researchers acknowledged that data centers in cities could produce two kinds of “heat island” impacts: namely, the combined effects of data heat and urban heat.

The study, released March 2026, estimated that the data heat island effect would impact more than 343 million people worldwide to face heat-related health risks, such as an increase in cardiovascular mortality.


This effect can accumulate if there are multiple data centers close to each other, the researchers saw a max increase of 9 degrees celsius, or +16.4°F.

## Try it Out

You can use AI Data Center Map (<https://aidatacentermap.org>) to estimate the Data Heat Island effect, which will calculate the projected temperature increase based on nearby hyperscalers.

In some areas of the United States, based on clusters of data center we also predict a major increase in heat up to the maximum +16.4°F seen in the study.





**When combined with the air pollutants and particulate matter produced by diesel generators, areas with high hyperscaler density would become a public health nightmare.**

## Noise Pollution

Aside from the risks that data centers pose to our water and our power supply, to our physical health and our financial health, there's the issue of their often-intolerable noise. Indeed, humans and animals living near data centers are likely to experience new, noticeable, and neuroses-inducing sources of sound pollution. In serious cases, constant sound pollution can lead to sleep deprivation and a host of health problems for neighboring households.

This is exactly what two New Jersey households experienced living within one mile of the DataOne USA facility, planned to be one of the largest data centers on the east coast.

In late May 2026, they filed a class-action lawsuit in federal court complaining about the constant industrial humming noise.

**“It’s a very loud industrial noise; sounds like a helicopter hovering, loud fans like you’re at an airport.”**

Gentile said in the complaint.

## MARA Cryptomine

A similar lawsuit was brought against a MARA Bitcoin in October 2024 by a group of Granbury, Texas residents. MARA (previously known as Marathon Digital) is the second-largest publicly traded holder of Bitcoin, expanding into Bitcoin mining and AI operations.



MARA Cryptomine

Although MARA Bitcoin is not a pure AI facility, crypto-mining facilities also use similar hardware and have comparable power demands.

The MARA Cryptomine began operating in 2022 with a planned capacity of 300 MW, powered by a gas plant next-door. By the fall of 2023, neighbors started suffering from the noise pollution. Residents described it as a loud vacuum cleaner running right next to them continuously, throughout the day and night, and vibrating throughout their entire house.

The noise got aggressive and unbearable when the company installed a wall around the facility.

The plaintiffs allege the crypto-mine has caused health issues ranging from migraines to hearing loss to sleep disturbances. The constant sound led to chronic issues of sleep deprivation, high blood pressure, anxiety, and general unease. Tinnitus and hearing loss were also common complaints.

MARA has denied the claims. It contends that the sound measurements taken at the site are below legal limits and that the facility is located in a well-established industrial zone.

One resident's dog began acting strangely, tearing its own fur out. Its anxious condition could not be treated with medicine, and eventually the pup needed to be put down. Most residents complained of dwindling numbers of wildlife in the surrounding areas.

Large data centers cause havoc to their neighbors in many different dimensions. Those who live nearby become trapped. They can't move — because who would want to trade places with them?

Yet another circumstance where residents are subjected to inhumane treatment by hyperscale data centers. Yet another circumstance where residents are subjected to inhumane treatment by hyperscale data centers.



16

# AI Baked In





**Even if one opposes the use of AI, these systems are now part of every major consumer technology platform. Whether it's speeding up your searches or erasing strangers out of your photos, AI is marketed as a convenient tool that's available when you need it.**

While AI may appear to be an add-on, the reality is that it is baked into modern platforms without an opt-out. Chances are, we are using AI without even knowing it.

It's in this way that millions of people who are opposed to data centers unknowingly justify the very existence of these same centers whenever they use modern big tech devices and services that entrap them unwittingly in deeply personal surveillance.

In this section, we examine the current state of embedded AI across device operating systems, web browser and search engines, and social media platforms.

## The AI New Normal

An analysis of big tech AI tools shows an almost identical approach to deployment. They arrive after a software update, are classified as system-level applications, or are even an integrated part of the OS itself. AI tools thus gain a level of trust that far surpasses other apps on our smart phones. This system-level classification allows AI tools direct access to the system in a way that is difficult to audit.

AI cannot be fully removed, though in some cases an option to “disable” is presented. As we’ve seen in the past through independent studies, simply turning off a feature doesn’t actually mean it’s “off.” The toggle switch may be there just to make you feel better. Tech companies hide the code from prying eyes. Their software and devices typically do whatever they want to in the background. Thus, the possibility always exists that your device is sharing data with AI services even after you disable it.

**Not one company has given users a way to completely remove AI from the product. It’s the new normal.**


## Big Tech AI Overview

This table gives an overview of every big tech product and its AI capabilities, making it easy for you to understand what AI operations are happening in your technology.

### Here's the explanation of each column in the table:

- 1. Company and Product:** identifying the company and the name of the product
- 2. Version:** indicates the software version or the year in which the AI features were introduced.
- 3. Default:** Whether the AI is on by default.  
(Y = Yes) (N = No)
- 4. Disable?:** Whether the user has the ability to toggle "off" features
  - a. Y\* =** You can toggle "off" features. But, due to the proprietary nature of the products mentioned, there is no way to confirm if the AI features are completely off.
  - b. N =** You cannot turn off all or some features.
- 5. Type:** The types of AI used in the product  
(Cloud or local AI)
- 6. Features:** What the AI is meant to do

| Company & Product        | Version                          | Default | Disable? | Type          | Features                                                                          |
|--------------------------|----------------------------------|---------|----------|---------------|-----------------------------------------------------------------------------------|
| Apple Intelligence iOS   | iOS 18.3, iPhone 15 Pro or Later | Y       | Y*       | Cloud & Local | Photo Analysis, Notification & Message Summaries, Writing Tools, Image Generation |
| Google Gemini Android    | Gemini                           | Y       | Y*       | Cloud & Local | Writing Tools, Image Editing                                                      |
| Samsung Galaxy AI        | S24 / 25                         | Y       | Y*       | Cloud & Local | Writing Tools, Summaries                                                          |
| Apple Intelligence macOS | Sequoia                          | Y       | Y*       | Cloud & Local | Writing Tools, Notification Summaries                                             |
| Microsoft Windows        | Windows 11, Late 2025            | Y       | N        | Cloud         | AI Assistant, "Memory"                                                            |
| Google Chrome            | Chrome 141, Oct 2025             | Y       | N        | Cloud         | Summarize, General AI Tasks                                                       |
| Google Search            | Nov 2024                         | Y       | N        | Cloud         | Summarize Search Results                                                          |
| Google Gmail             | Jan 2026                         | Y       | Y*       | Cloud         | Smart Compose, Reply, AI Writing                                                  |
| Meta Facebook            | May 2025                         | Y       | N        | Cloud         | Training on User Data, AI Assistant, Photo Scanning                               |
| X                        | August 2024                      | Y       | N        | Cloud         | Training on User Data, AI Assistant, Image & Video Generation                     |



17

# Settings You Should (Try To) Disable





# Phones

## iOS (Apple Intelligence)

Apple Intelligence is treated as part of iOS, is given elevated trust, and is allowed direct access to things like notifications, on-screen display, and even the phone's settings (this is how Siri can change iPhone settings for you).

Settings > Apple Intelligence & Siri > Disable

This is the main Apple Intelligence toggle which should disable features such as notification summaries, phone call summaries, and integrated writing tools.

Settings > Photos > Natural Language Search

This setting, turned on by default, scanned your entire photo library for people, and things—so you could search them naturally. It's unclear if the photo descriptions were sent to Apple.

## Settings > Photos > Scroll to Bottom > Enhanced Visual Search

Enhanced Visual Search was a feature enabled on all phones by default. It scanned your entire photo library for landmarks (buildings or notable objects) and sent a signature of these landmarks to Apple.

**This was enabled without so much as a notification or a blog post—much less your consent. It's unclear why Apple wants this data. The landmarks and buildings from photos across billions of devices along with locations could allow one to determine the location taken in any photo simply by analyzing the buildings or landmarks in the background.**

## Android (Gemini)

Google has given Gemini the highest levels of integration in the OS, incorporating the functionality inside the **Google app**, a special privileged system app that holds permissions that third-party apps can never obtain. From here it has direct access to system permissions and data, including calls and messages, microphone, location, contacts, and sensors.

In July 2025, the Google app also gained access to WhatsApp and system utilities. This occurred by default on a software update. The new Android 16 API indicates that system-level apps will receive even more abilities to accomplish additional tasks within the browser, calendar, camera, alarms and more.

**The only way to disable the AI is to disable the Google app itself:**

If you disable the Google app, many other things break as well. Voice search, assistant, discover feed, certain widgets, and Google services may behave unexpectedly.

Settings → Apps → Google → Disable

Settings → Apps → Gemini → Disable

*\*Note that the apps will be removed from your home screen and drawer.*

If you have problems, you can re-enable by:

Settings → Apps → Disabled Apps → Google App → Enable

Note that every system update will cause the Google app to re-enable itself. There is no permanent solution except to completely shift over to a de-googled device.

## Samsung

Samsung has confirmed that their Galaxy AI (starting with S24) is another system-level app that can only be disabled but not fully removed. It consists of a dozen features involving calls, photos, writing, notes, and more.

Settings → Galaxy AI → Disable individual features

Surprisingly, Samsung included a feature to avoid sending data off to the cloud. Turning it on will disable features such as generative edits and summaries.

Settings → Galaxy AI → Process data only on device

However, it's unclear if these settings will stay put after a software update.

# Computers

## Windows

After Microsoft's Copilot spread like an infection across Windows, from the file explorer to the humble Notepad, a recent blog post from Microsoft admits they went "too far."

Microsoft is now starting to remove Copilot from different apps, such as the Snipping Tool. In others, like Notepad, they simply removed the branding while keeping the AI features.

Before we give Microsoft any credit, though, critics note that they haven't **[fixed the changes that matter]**

- Windows cannot be used without a local account
- Ads are on the start menu and the lock screen
- It has system-wide telemetry
- OneDrive forcibly steals your data

To minimize its influence, you can attempt to disable CoPilot, but there's no guarantee it won't simply come back after the next update.



- Remove from Taskbar

Settings → Personalization → Taskbar — Turn off Copilot

## **Disable All AI Features using RemoveWindowsAI Script**

Warning: The RemoveWindowsAI Script will use advanced techniques to remove the majority of AI integrations. Antivirus software may falsely detect this script as malware.

Due to the tight integration of AI with Windows 11, it is possible for things to break. If this happens, simply run the script with the revert option.

We assume no liability for any errors or issues caused. You assume complete responsibility. Use at your own risk.

Watch the [instructional video on YouTube.](#)

If you run the script, there are some settings that cannot be turned off automatically.

- **Stop OneDrive from Scanning Faces In Your Photos**

Settings → Privacy and permissions → Features → People Section → Turn off

- **Disable Integrations with Apps**

Many apps have integrations with Copilot, such as Microsoft 365, and you may need to check at the app level.

Check for this at the app level:

File → Options → Copilot → Clear the Enable Copilot checkbox

# macOS

Apple Intelligence is also “baked in” to the macOS operating system and cannot be fully removed. However, there are toggles.

System Settings → Apple Intelligence & Siri  
→ Toggle Off → Confirm

This option should disable the AI features system-wide. **Users noticed that 5GB of disk space would clear up** after toggling it off, likely from the on-device AI models used by Apple Intelligence.

**Updates to your computer can also re-enable Apple Intelligence**, downloading GBs of AI models to your system.

Apple has made this option harder to find in recent updates, as if they’re trying desperately to prevent you from shutting it off.

# Browsers

## Chrome

Google's Chrome browser has AI features by default, starting on version 146, and your usage of the browser is used to train Google's AI.

There are a number of things you can do to disable Gemini:

In the URL bar, go to:

```
1 chrome://settings
```

### Gemini Sidebar Button

1. Search for "Gemini"
2. Toggle off Gemini in Chrome

This disables the Gemini button and keyboard shortcut.

## Page content sharing

1. Within chrome://settings
2. Search for “You and Google”
3. Click on AI Features → Page content sharing  
→ Turn Off

This disables Gemini from reading the content of your page and sending it for AI processing.

---

## AI Mode

In the URL bar, go to:

```
1 chrome://flags
```

- Search for “AI Mode Omnibox”
- Set to Disabled

The toggles above are the most accessible Chrome AI features. There are additional features that require a more technical approach.



# Social Media

## Meta

On January 1, 2025, Meta updated its Terms of Service to automatically grant it the right to use content from Facebook or Instagram.

### Permission to use content you create and share:

Specifically, when you share, post, or upload content that is covered by intellectual property rights on or in connection with our Products, you grant us a non-exclusive, transferable, sub-licensable, royalty-free, and worldwide license to host, use, distribute, modify, run, copy, publicly perform or display, translate, and create derivative works of your content.

Back in 2023, [Meta shared in a blog post that it was training its generative AI on public posts and content.](#)

The trend continues. Meta has tightly integrated AI into its mobile apps, with little way for you to adjust them. We've found a few options below:

## AI Mode

### Facebook & Instagram Apps

- Search Bar → Tap Meta AI Icon → Info → Mute → “Until I Change It”

### Disable Comment Summaries

- Settings & Privacy → Settings → Audience & Visibility → Posts -? Turn off comment summaries

You could also add Meta’s websites to your home screen to use the web version, which does not contain as much user-facing AI functionality.

However, this doesn’t prevent your content from being used by Meta. The only way to completely opt out is to delete your Meta account entirely.

## X (twitter)

**In January 2026, X updated its terms of service to claim rights to user content for any purpose.** X's AI system "Grok" is an indistinguishable part of the platform, which learns when you search, post, or even when you thumbs up or thumbs down posts.

Users are opted into letting Grok train on their information, but you can turn it off.

Settings and Privacy → Privacy and Safety →  
(scroll to) Data Sharing and Personalization → Grok  
→ Turn off "Allow your public data ... to be used for training"

## Conclusion


AI has made it possible for big tech companies to become acutely aware of the day-to-day lives of billions of people, opting them into AI systems without so much as a warning.

These AI systems are absorbing your social media, your personal messages, your personal emails, and your personal photos.

Not only is this practice justifying the use of massive data centers, but it is creating the largest and most sophisticated surveillance net the world has ever known.

The only way to fight back is to stand up for your rights. Use technology that respects you.

# YOUR DEVICE YOUR RULES




Privacy phones, laptops, and tablets that unlock new freedoms and protect your data.

[abovephone.com](https://abovephone.com)

 **above phone**



A high-contrast, black and white close-up of a fingerprint, showing the intricate ridges and valleys. The image is partially obscured by a red overlay on the right side.

18  
—

# Calculating A Model's Energy Costs



Our exploration thus far has been a high-level analysis of data centers and their power usage.

But how does that analysis magnify to a single person?

In this section, we'll look at the resource demands of chat-based LLMs and image/video generation.

## Calculating Energy Costs Of LLM Text Generation

In 2023, a new study at University of California Riverside captured everyone's attention. It estimated that between 20 and 50 ChatGPT queries evaporated 500 mL (17 oz) of fresh water—a normal-sized bottle of drinking water.

After this realization, Google and OpenAI were quick to put out their own estimates. Google's median Gemini query is said to use 0.24 Wh (watt hours).

However, successive estimates have reined this number in. EpochAI estimated a number about ten times less.

Based on a 500 word response using GPT-4o on modern hardware they arrived at .29 Wh, which can be compared to a few minutes of a smartphone charging.

How close is this number to an actual real-world setting when using Cloud AIs?

**When we use cloud AI platforms like OpenAI, Meta, and Anthropic, we're not using the open-source models that these researchers are using. We are missing lots of information:**

- **Models are opaque:** We don't know the true number of parameters for the proprietary models. OpenAI shared that their GPT-4 model had "hundreds of billions" of parameters without any further information. Because of this we can't calculate the power use accurately.
- **Multiple agents:** Instead of interacting with a single model, you are interacting with a complex AI service with a chain of agents processing your query behind the scenes
- **Huge clusters:** An AI data center may have any number of GPUs in the cluster, potentially far greater than ten—all of which need to be running for every query.

In essence, these calculations may be appropriate for open source models that you run on your machine, where your hardware is well-known and things can be verified.

But in the cloud, you are not interacting with a single model; you are interacting with an AI service.



This service involves routing your query through multiple layers of virtual AI agents: a routing agent that forwards your query, a guard agent that protects the system from malicious use, researchers that retrieve data from other sources, coders that develop assets and documents, and administrative agents that collect insights about people's usage.

Secondly, power usage is underestimated due to the rise of "memory" inside cloud AI agents, which remember the entire chat and even the user's entire history across some products.

This "memory" is fed back into each AI interaction as the "context," which then has to be reviewed by the AI on each query, requiring a significant amount of energy use. We can consider this "context" as input tokens to each query.

It is estimated that a long input of 10K tokens (several pages of information) would increase the cost per query to 2.5 Wh, and a very long input of 100K tokens (200 pages of text) would require 40 Wh of energy.

Due to the chains of multiple agents and long-term memory of cloud AI, we suggest the actual energy costs of cloud AI agents is several times higher than the 0.29 Wh per average query estimate, potentially three to five times on the low end—and more for advanced tasks with multiple agents.



19

# Guidelines For Economical AI Use



**There are two major areas of improvement that would minimize your AI footprint. Follow the guidelines below to minimize your energy costs.**

## **Use Local AI With Open-Sourced Models**

**Avoid Proprietary Models:** GPT-4o and Claude's Opus run in the cloud with little transparency. We don't know the number of parameters, the exact architecture details, or the weights. For example, all we know about ChatGPT's GPT-4o is that contains 'hundreds of billions' of parameters.

We also don't know which types of efficiency mechanisms GPT-4o or Opus use or how many GPUs are in a typical cluster, characteristics that all affect power usage.

**Use Open-Source models:** Mistral-7B, GPT-OSS, and DeepSeek v3 have published their parameter counts and architecture efficiencies. This means you're working with the real parameter count instead of an estimate. Additionally, running one of these models on your own hardware at home means you know exactly how much power it is using.

There are also projects like the AI Energy Score that calculate the energy required per 1K queries for text generation and image generation.

### See

<https://huggingface.co/spaces/AIEnergyScore/Leaderboard>



## Limit Your Tokens


Both cloud AIs and local AIs have different levers that can be tweaked to minimize your use of tokens and thus your power and water usage.

**System Prompt:** A system prompt gives the AI overarching instructions (“prompt engineering”) that tell the AI to be as concise as possible, unless you explicitly request more verbosity.

**Reasoning Effort:** Some AI models can be tuned to different levels of reasoning.


**Input Tokens:** Avoid uploading huge documents into the memory, and turn off memory when possible.

To replace usage of your cloud AI with local AI, watch our webinar that walks you through the setup: Empowerment with Ethical, Local, and Private AI



**Watch Webinar**





20

**The Purpose  
Behind Data  
Centers**





# One elephant remains in the room—in the form of a question: Why are so many data centers being built?

The answer: It depends whom you ask.

AI data center developers say the reason is to win the “AI cyberwar” against our enemies.

Governments say it’s a matter of national security. Yet, despite the “AI war” and despite national security, the common man or woman isn’t asking for more AI data centers. No one is complaining that their AI is too slow.

Could it be pure greed and profiteering that are impelling the growth of these gargantuan, natural-resource-draining structures?

We’ll explore a few theories.

## Speculative Bubble

First, let's get a sense of how much AI compute companies are currently using. To do this, we'll use operational data center announcements and investor reports from the AI companies.

When you size up all of Microsoft and Google's global data centers, with AI included, you see that each one has a total operational power of around 4 GW (see their environmental/sustainability reports).

These two companies are not only able to support billions of users across their standard services, they're also able to support large amounts of AI users. The latter include 400 million Microsoft 365 users, 650 million monthly Google Gemini users, and two billion users of Google AI Overview searches.

Keep in mind that this output is coming from a measly 4 GW!

Now, compare that 4 GW to the projected total power capacity of two new data center build-outs—17 GW for Project Matador in Texas and 9 GW for the Stratos Project in Utah—and it becomes apparent that their combined capacity of 26 GW will dwarf Google and Microsoft's global output several times over.

Besides Matador and Stratos, other new hyperscalers are just now coming online. Many of them are still in their initial stages of development. Fewer still have confirmed tenants.

## Take A Lesson From China

During 2023 and 2024, China announced that 500 new data center infrastructure projects would be constructed. Combined, these projects will drive up China's total power capacity to 32 GW, significantly less than the US' rough total of 50 GW capacity.

**Despite dealing with U.S export restrictions on its ability to build & train domestic AI models, China executed a nationalized AI plan and achieved an “AI-first” society:**

- AI has been deployed across the entire Chinese economy, notably in mid-sized and small business enterprises and in the education industry.
- Chinese cities integrate “smart” layers, such as AI, big data, IoT, and robotics, into their urban infrastructure to create an automated, data-driven governance system.
- Many AI chatbot services have adopted low-friction strategies such as free subscriptions

- Fully 600 million Chinese, or 42% of the country's population, are using AI regularly. In other words, there are more people using AI in China than are using the internet in the United States, and yet the US has built the majority of the world's AI compute.

The Chinese AI industry has achieved a lot, and yet by 2025 local outlets were reporting that **80% of new AI data centers in China were sitting completely idle.**

Investors looking for new projects after the COVID economic downturn jumped into AI, over-investing into building slews of data centers. They projected they could rent out all of the compute power to AI companies, but this failed for several reasons:

- Inexperienced hyperscale developers were off target when forecasting the type of infrastructure needed; they built data centers optimized for training instead of for inference.
- Overproduction of hyperscale resources led to excess supply, which pulled down rental prices.
- During the period when chip export controls were imposed by the US on China, Chinese AI companies were forced to innovate, leading to extremely efficient reasoning models like DeepSeek.

**This inactivity was a massive wake-up call, especially considering China's maturing adoption of AI. As a result:**

In April 2025, research institute Epoch AI estimated that between 2019 and March 2025 close to 75% of the world's AI compute was in the hands of the United States, with only 15% held by second-place China.

Given that China is struggling to put into use its much smaller proportion of data centers and given that it has a much larger, more AI-receptive population than does the United States, it's infeasible to think the U.S. will operate its data centers entirely on consumer & business demand.

The stark answer to this question is that many AI data center projects in the US are destined to fail. They cannot continue to count on growing consumer use, especially now that new innovations and consumer trends are shifting away from cloud-hosted models.

Our conclusion: The short-sighted, speculative greed of this industry will backfire, and there must be another reason why these data centers are being built.



## AI War Deterrence

When pressed, the main proponents of data centers in the US frame AI infrastructure development as a national security issue. They insist that building strong defensive capabilities such as massive data centers acts as a deterrent to an AI cyberwar by the US against foreign threats.

Infamously, Kevin O’Leary, developer of the Stratos data center in Utah, asked an interviewer if we in the West should put down our shovels and let China accelerate the pace of its data center construction.

O’Leary was alluding to concerns that China may use AI as a cyberweapon.

He shares the concerns of Dario Amodei, CEO of Anthropic, who has plenty of experience handling AI that could perform cyber-offensives. Anthropic was the victim of the first AI-orchestrated cyberespionage campaign, which required next-to-no human intervention.

In September 2025, foreign threat actors reportedly used Anthropic models as an integration layer, automatically running cyber-offensive campaigns based on a list of targets and given access to cybersecurity tools.

That incident may explain why Amodei alluded to AI being something like a nuclear weapon—and why he believes we are in a race to build the biggest and best weapons to deter and counter the AI-enabled threat actors (state-level actors and otherwise).

In this race to build a nuclear weapon, Anthropic released Claude Mythos, a highly advanced AI model, which it shared with only a few select organizations.

Mythos made headlines by generating successful exploits against target software in 72.4% of its trials. These trials were conducted almost autonomously.

Mythos was given a broad directive to find and exploit and was then able to experiment and confirm for itself.

Anthropic has been validating the Mythos exploits with expert human testers, followed by ethical reporting of the exploits to their authors. Its intention is to use Mythos and models like it for cyberdefense. Currently it is ramping up efforts to test and validate across open-source and proprietary groups alike.

**“The biggest risk AI creates is an enormous increase in vulnerabilities and the scale of hacking damage. . . . Ransomware attacks could expand to hit schools, hospitals, and banks.”**

Dario Amodei  
Speaking at an NY Finance Summit in April 2026

Amodei estimates the US has only six to twelve months to fix the vulnerabilities before these ransomware attacks happen.

Echoing his estimate is the National Institute of Standards and Technology’s (NIST) Center for AI Standard and Innovation. CAISI suggests China is only eight months behind frontier AI.

Last July, Secretary of State and Acting National Security Advisor Marco Rubio unveiled elements of “America’s AI Action Plan.”

**“Winning the AI race is non-negotiatble. America must continue to be the dominant force in AI to promote prosperity and protect our economic and national security.”**

Marco Rubio

The theme is a repeat of the Cold War-era Space Race. The message is that we need preemptive security and infrastructure against foreign AI threats.

Both China and the US are behaving as if there is an impending AI war, fighting both offense and defense simultaneously.

## The US government has a long history of economic policy to stifle China's AI manufacturing ambitions:

- In 2018, the US prevented the Netherlands from exporting Extreme Ultraviolet Lithography (EUV) machines to China (these are the machines that make advanced chips), and in 2023 this was followed by bans on less-sophisticated Deep Ultraviolet Lithography (DUV) machines.
- In 2022, the US banned the sale and export of high-performance AI chips to China across the board.

## China was able to counter:

- In 2025, China restricted exports to the US of rare earth elements that are necessary to build magnets, missiles, motors, and AI chips.
- China followed up with an even-broader wave of restrictions, adding additional more rare earth elements, products, and equipment.

Both countries de-escalated at the Busan Summit in South Korea, where they reached a basic consensus on tariffs and export controls and released their tight holds on chips and rare earth minerals.



But, though tensions relaxed on the surface, both countries are remaining focused on the defensive growth of AI infrastructure.

## **The US made major moves into AI and chip manufacturing:**

- The CHIPS Act, signed in 2022, allocated \$52.7 billion for domestic semiconductor manufacturing.
- The largest foreign investment ever made in the US was by Taiwan Semiconductor Manufacturing Company for a \$165 billion manufacturing plant in Phoenix, Arizona, along with fabrication and the first assembly plants in the US.
- Samsung, Intel, Tesla, and other major players announced new chip fabrication labs.

## **China continues to invest in its independence from the US-controlled chip manufacturing industry:**

- The Made In China initiative set a target of 70% domestic semiconductor fulfillment by 2025 (in 2026 it reached 30–35%), and Huawei's CEO, Ren Zhengfei, announced he is leading 2,000 companies in this effort.

- China is using alternatives to EUV, like the aforementioned Deep Ultraviolet (DUV), to produce less sophisticated chips at higher costs.
- Semiconductor Manufacturing International Corporation (SMIC), China's leading chip manufacturer, brought online a fabrication plant in 2025 and is in process to expand quickly across multiple sites.

Now that President Trump made a May 2026 in-person visit to China with a large delegation, including tech CEOs like Elon Musk and Jensen Huang, it appears on the surface, that both countries are attempting to smooth things over. Behind the scenes, though, they're still spending unfathomable amount of money arming themselves for an imagined AI war.

The publicly stated reason for each country wanting technological dominance is clear: preemptive security and infrastructure against "the other guy."

## But is that all there is to it?

## Robotics

Massive amounts of AI compute are needed for the development of robotics, although not where you might expect.

## Edge AI

When it comes to the day-to-day operations of a robot, AI algorithms and machine learning models don't rely on distant data centers and centralized cloud services for processing, instead they process AI tasks right within the robot's hardware itself, on a dedicated "edge" AI chip.

Edge AI is found on cameras, smart phones, autonomous vehicles, and robots that work on factory floors to warehouses and from hospitals and hotels to hazardous environments.

Having what is called "on-device intelligence" gives a robot instant decision-making capability. Built-in AI compute enables the robot to react quickly with only minimal delay.

## Self-Driving



A self-driving Apollo Go Robotaxi

The robot can thus make react instantly—avoiding obstacles, navigating roads, recognizing objects. Self-driving cars, for instance, use AI compute at the edge.

In-vehicle, purpose-built AI chips can react within milliseconds (or even sub-milliseconds) to road hazards. NVIDIA, Qualcomm, Rivian, and AutoX all produce Edge AI chips that power today's self-driving vehicles.

Advanced robotics applications use similar Edge AI chips whose low latency ensure safe and precise motor movements as well as providing extreme power efficiency to conserve energy.

## How To Train A Robot



Major robotics companies require the use of hyperscale data centers to train their robotics fleets. These companies often partner with NVIDIA, Microsoft, or OpenAI.

Humanoid robots are a new class of robots that are learning how to do everything humans can do.

When the humanoid robots are taught to simulate human body movements and do acrobatics, virtual training takes place by running millions of physics simulations.

When they are taught assist with chores, humanoid robots need a massive amount of data on the movements that human hands make when folding cloths or ironing or wiping tables or any number of repetitive tasks.



Investors paid \$6 billion in 2025 to launch new companies like Scale AI, Micro1, and Encord— AI training startups that pay gig workers to record themselves doing household chores (usually from a phone strapped to their head) in order to train humanoid robots to do the same.

After these workers, who are usually brought to the States from developing countries, film themselves performing chores endlessly, automated AIs review the videos. Then humans label the videos with more information.

Meanwhile, in China, local governments have already invested in 40 state-owned robot data collection centers, where humans train robots in folding laundry, stocking shelves, and cleaning tables.

Whether it's running millions of physics simulations, analyzing hundreds of thousands of videos, or generating artificial training videos, AI data centers are used in every part of the pipeline.

## Robot Training Factories

The hundreds of thousands of hours of “chores” footage are eventually batched together and fed into large AI training pipelines. Robotic companies primarily partner with cloud AI operators for this effort.

Tesla is planning its first humanoid robot, Optimus, using 100,000 GPUs at its Cortex 1 gigascaler in Texas, shrinking down training pipelines to hours instead of weeks.

One notable robotic pioneer, Boston Dynamics, has a longstanding partnership with Google and trains its robots on Google Cloud data centers.

Smaller companies, among them Agility Robotics and Germany’s NEURA Robotics, both of which specialize in humanoid robotics, are backed by Amazon Web Services (AWS) data centers.

Amazon is expanding its use of humanoid robots in fulfillment centers and is actively testing them for last-mile delivery.

## Surveillance And Predictive Policing

With automated surveillance systems like Flock cameras on the rise in the United States, one must wonder if AI data centers could be used to usher in a surveillance state never before conceived by science fiction.

In this envisioned future, all video cameras would feed into a central AI hub, identifying people and their actions in real time—24/7.

A system similar to what we're describing already exists in mainland China. The aptly named "Skynet" system reportedly monitors and pools over 700 million camera feeds throughout China, linking roads, residential areas, and public areas into a single database. A sister system, dubbed "Sharp Eyes," extends the vision out to remote villages and townships.

These systems are integrated with facial recognition so that authorities can instantly identify individuals, determine their legal status, track their movements throughout the network.

In major Chinese cities, where there is roughly one camera for every 1.6–2.0 residents, it's exceedingly difficult to avoid being filmed.

Yet, despite the proliferation of cameras, ordinary citizens who ask for help rarely receive assistance. Missing children are never found and hit-and-runs remain unsolved, but speeding tickets arrive on schedule.

Even though the government claims the cameras are monitoring and recording in real-time, we think it unlikely that Skynet is real-time.

Why? Because 700 million simultaneous video streams would require petabits of network bandwidth—far beyond the capability of any country.

The Oxford Institute for Technology and Justice suggests that Skynet is primarily used for targeting purposes—that is, for picking out specific individuals. These individuals are usually chosen from other AI systems, which have spotted, flagged, and deemed their behavior “suspicious.” What constitutes “suspicious” could range from the use of a VPN to the purchase of some frowned-upon product or service.

Once a target has been identified, Skynet can be used to find where the individual is. It looks up their faceprint on stored camera feeds. The underlying facial recognition software is Megvii’s Face++, which is a black box proprietary service that can detect 106 facial landmarks within a millisecond.

So, although China's entire surveillance camera network is not real-time, it is next-to-real-time for a smaller group of targets. Many things remain unclear such as the amount of targets it can track, and the total size of the AI infrastructure required to run the system.

In the United States, a similar system has been rolled out for vehicles. The 80,000+ Flock cameras throughout the country recognize vehicles, license plates, and special features.

Flock is a near-real-time system that can track targets across the network.

**Typically these surveillance systems use a combination of Edge AI (to recognize faces and features quickly) and a cloud AI integration layer (which does further learning and coordination).**



## China also uses extensive AI systems for policing:

- **Integrated Joint Operations Platform — IJOP** is a data fusion tool that automatically flags individuals for suspicion. It tracks where phones, ID cards, and vehicles go, as well as electricity use, internet usage, and more. IJOP is integrated with data doors; these are physical gateways people walk through that automatically identify them by retrieving identifiers from their phones. If their identity is flagged, officers manning the data door can be notified to initiate automated detainment and interrogation.
- AI police teams deploy hundreds of AI models for fraud detection, crime prevention, case investigation, automating evidence gathering, and even sentencing for crimes.

Altogether, AI infrastructure powers the backbone of China's far-reaching surveillance state. As other countries look to provide data-fusion approaches to law enforcement, AI infrastructure may increasingly be used to combine data from financial records, live camera feeds, smart infrastructure, and control gates.

**This is a fearsome vision for the future—but one made possible by hyperscale data centers.**



# 21 - Solutions





We're moving through a dark place where multiple crises compete for our attention, yet we feel the AI ecology issue is worth giving our devoted attention. It marks the 'peak' of civilization, breaking records in every dimension in history; dollars invested, buildings expanded, power & water used, and pollution generated. This issue impacts all living beings.

We're unclear what the future will look like, but we must act against the growing drain on our natural resources. This solutions section summarizes the most important actions to undertake.

## Set Down Your Roots

Trees integrate with their environment through their roots, measuring the amount of moisture and nutrients in the soil — and communicating with other plants through the mycellium network.

We must also work to establish ourselves in our communities, make people aware of the larger issue, and form robust communication networks. Join local groups, and speak up about what you've learned.

Although data centers are now entering the mainstream, few have a complete understanding or education on the topic.

By reading this report, you have something incredibly important to share that affects everyone in your town.

## Try out the following activities:

- Research your local data center using the process in this report
- Hold a small get-together, discussion circle, or dinner to discuss the issue
- Speak at your county or city council hearing.
- Understand the legal procedure to approve hyperscalers in your area

## Win Over Your Neighbors & Officials: People Power

Ultimately, this is a people issue. Local residents approve zoning changes, county officials approve tax abatements, and construction workers build the facilities.

People issues are solved with people power.

From our conversations with activists around the country, we realized that data center developers are connecting with local decision-makers before residents get a chance.

Executives will visit city and county officials and pitch them on the idea of having a data center in their town. Oftentimes, they'll drive them out to a nearby data center and impress upon them the "benefits" of building a hyperscaler locally.

Residents aren't connecting with their local officials to share the other side of the story, and should make the effort. If we wait until a rezoning hearing, its far too late. Two minutes of discussion before a vote may not change anything. We must build the connection far before the hearings, connecting on a personal level, outside of official settings.

## Here are some activities you might try:

- Having direct person to person conversations with every member of city council / county commissioners office.
- Holding a lunch or open discussion with residents in the area and invite your councilmembers.
- Sharing testimonials from residents who live next to a data center
- Planning a tour to show the heat and environmental impacts of nearby data centers



## Minimize Your Use of Cloud AI

As we've covered in this report, Cloud AI is directly tied to power usage, water usage, and the bottom line of AI companies.

Not only does paying for AI subscriptions perpetuate the use of AI data centers, but the information you share with Cloud AI is used to further refine and tweak Cloud AI models.

For these reasons, we recommend people to minimize their use of Cloud AI and switch to Local AI that can run directly on your computer.

Local AI is powered by open-source models that are often initially trained on AI data centers. However, once the training period is complete and the model is released to the public, it has no dependency or use of the AI data center center.

This means it runs completely off the power of your device, and doesn't require an internet connection. We predict that the future will replace our reliance on Cloud AI services to local AI services running in your home.

## Switch to Sovereign Devices

Our discussion of solutions would not be complete without discussing phones and laptops. We currently live in a world where 99% of people use devices that are controlled by the companies that develop the software. Forced updates are rolled out on a regular basis to billions of people.

These forced updates permanently change the operating system you use day-to-day; they bring about features like digital wallets (Digital ID), health surveillance (Contract Tracing), and of course — AI.

In the ‘AI Baked In’ section of this report, we discussed how these changes fundamentally change the software on your device, they can’t be isolated or removed, and even if you manage to disable them — you risk their return with the next update.

AI on your phone or laptop is especially worrisome due to how it can integrate with your entire device. The AI layer on modern devices are capable of much more than polishing your emails, the possibility remains they can learn from all your activity across apps, messaging, and more.

**Many people are waking up to the realization they replace their phone or laptop with an open-source alternative where they remain in control of their device, not a third party company.**

**These devices offer all of the same apps and programs that normal phones do, but come without the extractive software that we're all increasingly growing tired of.**

**No advertising, no digital ID, and most importantly — no forced AI. We envision a future where billions are using these private devices to replace their dependence on abusive software ecosystems.**

# Reclaim your phone & laptop.

## Ready to explore a de-googled phone or a Linux laptop?


- Watch our **Rise Above Ecosystem Overview** to learn about the phone, laptop, and how they protect you
- Head over to **Above Phone** to browse the current deals and to chat with our team

## Features Include:

- Call & text with any cell service
- Download popular apps privately
- Discover friendly alternatives
- Private web browsing & search
- Ad-Free Videos, Music, Podcasts
- Encrypted messaging, voice, and video calls
- Private WiFi, Bluetooth, Hotspot
- Offline maps & navigation
- Create isolated containers for apps

## Interested in running Local AI?

Get the scoop and watch our Empowerment with Ethical AI webinar. It walks you through step-by-step how to setup your own useful local AI system that can be utilized as a private translation service, document summarizer, and so much more — all free and local to your machine.



**Watch Webinar**

<https://learn.abovephone.com/course/empowerment-with-ethical-local-and-private-ai/>



22

-

# Conclusion: A Spiritual Tipping Point



The theories presented above are more than plausible. What threatens to make them possible is the greed and shortsightedness of industry leaders and the obsessive insecurity and paranoia of nation states.

Even our theories remain incomplete due to the sheer scale of the planned AI buildout. Our estimates with the AI Data Center Map indicate that 200GW of AI compute could be built by 2030. That's 40% of the current U.S power supply.

At a high-level, this power level is far greater than what is required for consumer AI, cyber-defense, or even a comprehensive surveillance system. It still remains unclear why this scale of compute is needed.

But perhaps the hyperscalers serve as a wake-up call for all people, a gentle nudge to pay more attention to the inner-workings of our planet. It marks a spiritual tipping point, where we turn inwards, seeking understanding, and expressing gratitude for our planet.

The decision-makers behind AI may be operating out of ignorance, not knowing the true costs of their actions. Although the hyperscalers that are built today will likely be financially insolvent and defunct within five years, their existence will have come at a major cost to our environment.

The extent of that cost is up to us. Things may get worse before they get better. Sometimes we need to learn the hard way, the collective realization that we're all equally responsible for the planet we call "home."

The decision-makers must also find this awareness. If hyperscalers destroy our natural landscape and make food production difficult, they also must live with the mess that they helped create.

Thankfully, we're starting to see the tide turn. Gradually, more and more of us are waking up to the grim realities of AI data centers and the despotic surveillance network they enable.

You, my dear reader, are a crucial piece in this puzzle. Will you agree to make hard choices instead of opting for convenience and comfort? Will you continue learning the cold, hard facts? Will you resist, especially at the local level, where pushback counts most? And will you inform those around you—perhaps by sharing this report with them?

Today and every day, please join me in living life consciously and with awareness. Together, we can pray, aligning our actions with peace, fulfillment, and happiness for all life on this planet.

# Acknowledgements

To my mother, who nurtured my  
independent thinking.

To my mentors who inspired  
greater visions.

To my friends and team, who have  
supported me every step of the way.

To the earth, which carries me in her  
loving embrace.

To the source from which all life springs.

I am endlessly grateful.

# About The Author



## Hakeem Anwar

A conscious technologist building  
private & friendly solutions.

Since 2021, as the founder of #TakeBackOurTech  
(<https://takebackourtech.org>), Hakeem guides thousands away  
from abusive software ecosystems that dominate the internet, giving  
them solutions in his in-depth guides, articles, and weekly show.

He also serves as founder & CEO of **Above**  
(<https://abovephone.com>), where he translates that same vision  
into tangible products; private phones, laptops, and internet utilities  
made easy to use.

Most importantly, he encourages his audience to tap into their  
organic intelligence through natural practices such as meditation,  
breathwork, and yoga.



# Contact

Do you have questions,  
feedback, or a message for  
Hakeem or the team?

**Contact us Here**

# AI Statement

As part of responsible AI usage, we have included the areas where AI is used, and where it isn't.

**Research** - The process of finding sources relevant to the report, finding useful research and information. The research for this report was primarily Human-Only, with Cloud AI used for assistance in the following sections:

- Who's Driving These Data Centers?
- The Purpose Behind Data Centers
- How To Spot a Data Center
- Gigascale Projects Transforming the U.S

**Writing & Editing** - Human-Only, Hakeem Anwar & Editor

**Graphics** - Human-Only, Above Design Team

**Citations** - AI-Generated

23  
-

# Resources



# Reports & Papers

## Peer Reviewed Articles

Jegham, M. Abdelatti, L. Elmoubarki, and A. Hendawi, "How hungry is AI? benchmarking energy, water, and carbon footprint of LLM inference," arXiv preprint arXiv:2505.09598, 2025.

Arzumanyan, M., Rodriguez Calzado, E., Lin, N., Bahadur, V., Das, J., Ko, T. L., & Koesterke, L. (2025). Geospatial suitability analysis for data center placement: A case study in Texas, USA. *Sustainable Cities and Society*, 131, 106687. <https://doi.org/10.1016/j.scs.2025.106687>

Chen, X., Wang, X., Colacelli, A., Lee, M., & Xie, L. (n.d.). Electricity demand and grid impacts of AI data centers: Challenges and prospects. Texas A&M University Consortium on AI and Large Flexible Load (CALL).

Gour, N., Ortiz, L., & Maibach, E. (2026). Health implications of the rapid rise of data centers in Virginia: An exploratory assessment. *Frontiers in Climate*, 8, 1648912. <https://doi.org/10.3389/fclim.2026.1648912>

Jasechko, S., Seybold, H., Perrone, D., Fan, Y., Shamsudduha, M., Taylor, R. G., Fallatah, O., & Kirchner, J. W. (2024). Rapid groundwater decline and some cases of recovery in aquifers globally. *Nature*, 625, 715–721. <https://doi.org/10.1038/s41586-023-06879-8>

Jegham, N., Abdelatti, M., Koh, C. Y., Elmoubarki, L., & Hendawi, A. (2025). How hungry is AI? Benchmarking energy, water, and carbon footprint of LLM inference (arXiv:2505.09598v6) [Preprint]. arXiv. <https://arxiv.org/abs/2505.09598>

Lei, N., Lu, J., Shehabi, A., & Masanet, E. (2025). The water use of data center workloads: A review and assessment of key determinants. *Resources, Conservation and Recycling*, 219, 108310. <https://doi.org/10.1016/j.resconrec.2025.108310>

Li, P., Yang, J., Islam, M. A., & Ren, S. (2025). Making AI less "thirsty": Uncovering and addressing the secret water footprint of AI models (arXiv:2304.03271v5) [Preprint]. arXiv. <https://arxiv.org/abs/2304.03271>



# Government Reports

International Energy Agency. (2025). Energy and AI (World Energy Outlook Special Report). IEA. <https://www.iea.org/>

Morey, M. (2024, October 1). Data center owners turn to nuclear as potential electricity source. U.S. Energy Information Administration, Today in Energy. <https://www.eia.gov/todayinenergy/detail.php?id=63304>

Shehabi, A., Smith, S. J., Horner, N., Azevedo, I., Brown, R., Koomey, J., Masanet, E., Sartor, D., Herrlin, M., & Lintner, W. (2016). United States data center energy usage report (LBNL-1005775). Lawrence Berkeley National Laboratory.

Shehabi, A., Smith, S. J., Hubbard, A., Newkirk, A., Lei, N., Siddik, M. A. B., Holecek, B., Koomey, J., Masanet, E., & Sartor, D. (2024). 2024 United States data center energy usage report (LBNL-2001637). Lawrence Berkeley National Laboratory.

U.S. Geological Survey. (2003, November). Ground-water depletion across the nation (USGS Fact Sheet 103-03). U.S. Department of the Interior.

# Academia Research

COMPASS Research Consortium. (2025, December 11). Water use requirements for data centers in Texas: A white paper on the evolving demands of water use in data center infrastructure in Texas. Bureau of Economic Geology, The University of Texas at Austin. <https://repositories.lib.utexas.edu/items/454bb243-7b0b-4453-8b28-1a47448c78dd>

Eckstein, G., & Hardberger, A. (Eds.). (2020, April 1). Groundwater laws and regulations: A preliminary survey of thirteen U.S. states, Vol. I (2nd ed.). Texas A&M University School of Law & St. Mary's University School of Law. <https://doi.org/10.37419/EENRS>.  
USStateGroundwaterLaws.2020

Eckstein, G., & Hardberger, A. (Eds.). (2022, April 15). Groundwater laws and regulations: Survey of sixteen U.S. states, Vol. II. Texas A&M University School of Law & Texas Tech University School of Law. <https://doi.org/10.37419/EENRS.USStateGroundwaterLaws.2022>

Krishnamoorti, R., & Datta, A. (2025, Summer). The future of the electric grid in Texas: Opportunities and challenges in the next decade [Research white paper]. UH Energy, University of Houston.

Spenrath, M., & Datta, A. (2025, May). Powering Texas' digital economy: Data centers and the future of the grid — Part One. Houston Advanced Research Center (HARC) & University of Houston.

<https://harcresearch.org/wp-content/uploads/2025/05/Powering-Texas-Digital-Economy-Part-One-Final-1.pdf>

Spenrath, M., & Datta, A. (2025, May). Powering Texas' digital economy: Data centers and the future of the grid — Part Two. Houston Advanced Research Center (HARC) & University of Houston. <https://harcresearch.org/wp-content/uploads/2025/05/Part-2.-Powering-Texas-Digital-Economy-Final.pdf>

## Sustainability Reports

Amazon Web Services. (2024). 2024 Amazon sustainability report: AWS summary. Amazon.

Bizo, D., Ascierto, R., Lawrence, A., & Davis, J. (2021, September 1). Uptime Institute global data center survey 2021: Growth stretches an evolving sector (UI Intelligence Report 51). Uptime Institute.

Google. (2025). Google environmental report 2025. Google LLC.

Meta Platforms, Inc. (2024). 2024 sustainability report: For a better reality. Meta.

Meta Platforms, Inc. (2025). 2025 sustainability report: Connecting to a better reality [Reporting on fiscal year 2024]. Meta.

Microsoft Corporation. (2025). 2025 environmental sustainability report: Accelerating progress to 2030 [Reporting on 2024 fiscal year]. Microsoft.

## Trade Analysis

Ahmad, R. (2024, March 4). Engineers often need a lot of water to keep data centers cool. Civil Engineering Magazine, American Society of Civil Engineers (ASCE).

<https://www.asce.org/publications-and-news/civil-engineering-source/civil-engineering-magazine/issues/magazine-issue/article/2024/03/engineers-often-need-a-lot-of-water-to-keep-data-centers-cool>

Jennings, C. (2025, July 23). The cloud is drying our rivers: Water usage of AI data centers. Ethical CEO.

You, J. (2025, February 7). How much energy does ChatGPT use? Epoch AI Gradient Updates. <https://epoch.ai/gradient-updates/how-much-energy-does-chatgpt-use>

## Activist Groups

MediaJustice. (2025, September 17). The people say no: Resisting data centers toolkit. MediaJustice. <https://mediajustice.org/resource/the-people-say-no-toolkit/>

Data Center Action Coalition. (n.d.). We won't let AI run us dry: No to data centers in San Marcos, TX [Community pamphlet]. DCAC, San Marcos, Texas.



# Web Links

## Who Is Demanding Data Centers?

<https://techcrunch.com/2026/02/28/billion-dollar-infrastructure-deals-ai-boom-data-centers-openai-oracle-nvidia-microsoft-google-meta/>

<https://techcrunch.com/2026/02/28/billion-dollar-infrastructure-deals-ai-boom-data-centers-openai-oracle-nvidia-microsoft-google-meta/>

<https://futuresgroup.com/insights/ai-capex-2026-the-690b-infrastructure-sprint/>

<https://techcrunch.com/2026/02/28/billion-dollar-infrastructure-deals-ai-boom-data-centers-openai-oracle-nvidia-microsoft-google-meta/>

<https://futuresgroup.com/insights/ai-capex-2026-the-690b-infrastructure-sprint/>

<https://openai.com/index/announcing-the-stargate-project/>

# Gigascale Project Impacts

<https://www.insideindianabusiness.com/articles/state-investigates-groundwater-issues-near-amazon-gm-projects-in-st-joseph-county>

<https://archive.is/AmzUL>

<https://www.businessinsider.com/indiana-big-techs-ai-data-centers-energy-amazon-google-microsoft-2025-2?op=1>

<https://goodjobsfirst.org/data-center-tax-abatements-why-states-and-localities-must-disclose-these-soaring-revenue-losses/><https://archive.is/ICOuF>

<https://www.wthr.com/article/news/local/13-investigates-you-pay-sales-tax-some-indiana-data-centers-didnt-heres-what-we-found/531-ae54b7a8-d19c-47ae-9375-d9edf4e4495f>

<https://www.glennklockwood.com/garden/systems/Rainier>

<https://truthout.org/articles/big-tech-data-centers-compound-decades-of-environmental-racism-in-the-south/>

<https://www.selc.org/press-release/groups-threaten-lawsuit-over-xais-unpermitted-gas-turbines-in-mississippi/>

<https://newsletter.semianalysis.com/p/xais-colossus-2-first-gigawatt-datacenter>

<https://www.bloomberg.com/news/articles/2025-12-30/musk-s-xai-to-expand-colossus-data-center-information-reports>

<https://www.selc.org/news/resistance-against-elon-musks-xai-facility-in-south-memphis-gets-stronger/>

<https://www.selc.org/news/experts-say-data-center-growth-is-speculative/>

<https://archive.is/8bnjZ>

<https://mississippitoday.org/2025/11/24/southaven-residents-fear-pollution-complain-of-noise-from-elon-musks-xai-data-center-turbines/>

<https://time.com/7308925/elon-musk-memphis-ai-data-center/>

<https://www.datacenterdynamics.com/en/news/xai-importing-power-plant-from-abroad-for-new-memphis-data-center-elon-musk-claims/>

<https://naacp.org/articles/elon-musks-xai-threatened-lawsuit-over-air-pollution-memphis-data-center-filed-behalf>

<https://www.commercialappeal.com/story/money/business/development/2025/07/15/elon-musk-xai-in-memphis-colossus-2/85215548007/>

<https://datacenters.atmeta.com/richland-parish-data-center/>

<https://earthjustice.org/press/2026/consumer-groups-alarmed-as-louisiana-psc-declines-to-take-up-probe-into-meta-risky-financing-deal>

<https://sherwood.news/tech/hyperion/>

<https://www.all4energy.org/watchdog/meta-data-center-to-cause-entergy-bill-increase/>

<https://www.wbrz.com/news/meta-s-louisiana-data-center-causing-concern-over-water-usage>

[https://web.archive.org/web/20260215184849/https://www.theadvocate.com/baton\\_rouge/news/environment/meta-louisiana-data-center-water-ai/article\\_02443429-0857-5574-a194-64839996bc52.html](https://web.archive.org/web/20260215184849/https://www.theadvocate.com/baton_rouge/news/environment/meta-louisiana-data-center-water-ai/article_02443429-0857-5574-a194-64839996bc52.html)

<https://web.archive.org/web/20250715015858/https://www.nytimes.com/2025/07/14/technology/meta-data-center-water.html>

<https://www.sierraclub.org/sierra/rural-louisiana-meta-s-new-data-center-promises-growth-what-cost>

<https://pubs.usgs.gov/publication/sir20225027/full>

<https://www.hrw.org/news/2024/01/25/us-louisianas-cancer-alley>

<https://thecoolingreport.com/intel/how-direct-to-chip-liquid-cooling-works-engineering-guide.html>

<https://datacenters.atmeta.com/richland-parish-data-center/>

<https://fortune.com/2025/08/24/meta-data-center-rural-louisiana-framework-ai-power-boom/>

<https://fortune.com/2026/02/04/meta-hyperion-ai-data-center-louisiana-expansion/>

<https://www.all4energy.org/watchdog/meta-data-center-to-cause-entergy-bill-increase/>

<https://earthjustice.org/press/2026/louisiana-regulators-asked-to-investigate-shady-meta-financing-deal-that-could-leave-households-paying-for-tech-giants-electricity>

<https://earthjustice.org/press/2026/consumer-groups-alarmed-as-louisiana-psc-declines-to-take-up-probe-into-meta-risky-financing-deal>

<https://sherwood.news/tech/hyperion/>

[https://www.nola.com/news/environment/meta-data-center-ai-louisiana-entergy/article\\_662ae263-22e1-4346-823c-d9a2ab498f4c.html](https://www.nola.com/news/environment/meta-data-center-ai-louisiana-entergy/article_662ae263-22e1-4346-823c-d9a2ab498f4c.html)

<https://fortune.com/2025/09/23/openai-oracle-and-softbank-expand-stargate-with-five-new-ai-data-centers-staging-texas-media-spectacle-to-counter-critics/>



<https://www.whitehouse.gov/>

<https://www.distilled.earth/p/these-data-centers-are-getting-really>

<https://www.datacenterwatch.org/newsletter>

<https://techcrunch.com/2025/12/08/environmental-groups-call-for-halt-to-new-data-center-construction/>

## Our Water Systems

<https://www.edwardsaquifer.net/intro.html>

## Water Rights

<https://nationalaglawcenter.org/overview/water-law/>

[https://scholarship.law.tamu.edu/cgi/viewcontent.  
cgi?article=1011&context=nrs-publications](https://scholarship.law.tamu.edu/cgi/viewcontent.cgi?article=1011&context=nrs-publications)

[https://web.archive.org/web/20240908015619/https://twri.tamu.  
edu/publications/txh2o/2024/summer-2024/understanding-  
groundwater-law/](https://web.archive.org/web/20240908015619/https://twri.tamu.edu/publications/txh2o/2024/summer-2024/understanding-groundwater-law/)

[https://amarillotribune.org/2025/11/05/just-facts-fermi-america-  
and-project-matador/](https://amarillotribune.org/2025/11/05/just-facts-fermi-america-and-project-matador/)

# Data Center Project Lifecycle

<https://www.arcgis.com/home/item.html?id=d77284d1a97a426eaf960bab4764999c>

<https://comptroller.texas.gov/taxes/property-tax/county-directory/bosque.php>

<https://comptroller.texas.gov/economy/development/prop-tax/ch312/>

<https://comptroller.texas.gov/economy/development/search-tools/ch312/abatement-zone-report-details.php?id=000016554>

<https://comptroller.texas.gov/economy/development/search-tools/ch312/abatements-simple.php>

<https://comptroller.texas.gov/economy/development/search-tools/ch312/abatements-details.php?id=000004667>

<https://comptroller.texas.gov/economy/development/search-tools/ch312/abatements-simple.php>

<https://goodjobsfirst.org/cloudy-with-a-loss-of-spending-control-how-data-centers-are-endangering-state-budgets/>

<https://www.cnn.com/2025/06/20/tax-breaks-for-tech-giants-data-centers-mean-less-income-for-states.html>

<https://goodjobsfirst.org/data-center-moratorium-bills-are-spreading-in-2026/3>

## Why

<https://www.technologyreview.com/2026/04/01/1134863/humanoid-data-training-gig-economy-2026-breakthrough-technology/>

<https://www.technologyreview.com/2025/03/26/1113802/china-ai-data-centers-unused/>

<https://www.technologyreview.com/2025/03/26/1113802/china-ai-data-centers-unused/>

<https://www.unite.ai/us-sanctions-backfire-huaweis-ai-chips-accelerate-chinas-self-reliance/>

<https://apnews.com/article/huawei-ascend-ai-chips-nvidia-superpod-1835ff00671858955f482f10122600f2>

<https://www.digitimes.com/news/a20251119PD234/huawei-investment-supply-chain-technology-equipment.html>

<https://epoch.ai/data-insights/ai-supercomputers-performance-share-by-country>

<https://www.brookings.edu/articles/how-china-harnesses-data-fusion-to-make-sense-of-surveillance-data/>

<https://en.people.cn/n3/2026/0303/c90000-20430758.html>

<https://www.csis.org/analysis/deepseek-huawei-export-controls-and-future-us-china-ai-race>

<https://www.chinadaily.com.cn/a/202410/23/WS6718580ca310f1265a1c9146.html>

<https://www.jll.com/en-us/insights/market-dynamics/north-america-data-centers>

<https://www.visiontimes.com/2026/04/12/china-surveillance-state-800-million-cameras-total-monitoring-and-public-backlash.html>

<https://www-cdn.anthropic.com/08ab9158070959f88f296514c21b7facce6f52bc.pdf>

<https://restofworld.org/2026/china-robots-training-centers-workers/>

<https://eu.36kr.com/en/p/3709346127933831>

<https://www.goldmansachs.com/insights/articles/chinas-ai-providers-expected-to-invest-70-billion-dollars-in-data-centers-amid-overseas-expansion>

<https://www.prnewswire.com/news-releases/ubtech-humanoid-robot-walker-s2-begins-mass-production-and-delivery-with-orders-exceeding-800-million-yuan-302616924.html>

<https://www.technologyreview.com/2025/03/26/1113802/china-ai-data-centers-unused/>

<https://techblog.comsoc.org/2026/02/16/china-vs-u-s-generating-power-for-ai-data-centers-as-demand-soars/>



# Health Impacts

<https://fortune.com/2026/04/01/ai-data-centers-heat-island-hyperscalers/>

<https://www.frontiersin.org/journals/climate/articles/10.3389/fclim.2026.1648912/full>

<https://wilder.vcu.edu/news-and-events/news-posts/northern-virginia-data-center-air-pollution-rivals-power-plant-emissions-vcu-research-finds.html>

<https://news.vcu.edu/article/northern-virginia-data-center-air-pollution-rivals-power-plant-emissions>

[https://scholarscompass.vcu.edu/cgi/viewcontent.cgi?article=1003&context=isee\\_pubs/](https://scholarscompass.vcu.edu/cgi/viewcontent.cgi?article=1003&context=isee_pubs/)

[https://scholarscompass.vcu.edu/cgi/viewcontent.cgi?article=1003&context=isee\\_pubs/](https://scholarscompass.vcu.edu/cgi/viewcontent.cgi?article=1003&context=isee_pubs/)

[https://scholarscompass.vcu.edu/isee\\_pubs/3/](https://scholarscompass.vcu.edu/isee_pubs/3/)

<https://www.frontiersin.org/journals/climate/articles/10.3389/fclim.2026.1648912/full>

**We appreciate  
you taking the  
time to review this  
report and hope  
it provides useful  
insight going  
forward.**



© 2026 Above LLC

**Disclaimer:** This report is for informational purposes only and reflects the author's opinions. It is not investment, financial, legal, or health advice, and we accept no liability for reliance on its contents. **Trademark Notice:** All trademarks are the property of their respective owners; mention does not imply affiliation or endorsement.