THE DATA CENTER NEXT DOOR

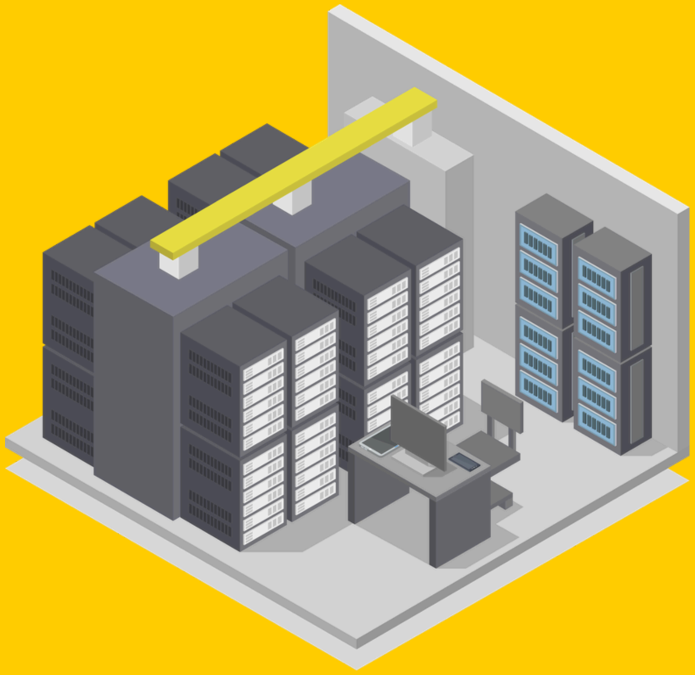

Key Considerations for Communities Navigating Data Center Development

USC
Annenberg
Center for Climate Journalism and Communication

APRIL 2026

Produced in conjunction with Season 3 of Electric Futures

TABLE OF CONTENTS

WHAT IS A DATA CENTER?

A data center is a dedicated physical facility – a room, building, or group of buildings – that houses an organization's critical IT infrastructure. It serves as a centralized hub for all the computing and networking equipment to run the digital aspects of the business.

In previous eras, a data center was a small room with servers and switching equipment. The rise of cloud computing led to a demand for storage for vast amounts of data that are generated through the use of the camera apps on our phones, text and video data developed for YouTube, WhatsApp, TikTok, and beyond.

In addition, corporations have found that use of cloud applications, such as Microsoft Office, Concur, Workforce, or Oracle are efficient and effective, leading to a shift in their digital operations from computers housed in a company to "the cloud." The cloud consists of large computers with large amounts of storage that are heavily protected from cyber threats.

The compound annual growth rate for cloud computing is about 25%. Thus, after 5 years, the amount of cloud computing increases by a factor of three, increasing the number of large-scale computing facilities needed by the same factor. Computation in the cloud includes activities like:

- Infrastructure as a service (outsourcing the maintenance of computers and storage devices),
- Platform as a service (developing, running and managing software)
- Software as a service (using applications like SAP for financial and HR functions)

This is big business. Cloud computing centers of Microsoft, Google, Amazon and Meta – companies referred to as hyperscalers – routinely cover millions of square feet. And the demand for these services simply grows. This growth – giving us smartphones, search engines, social media and enormously powerful software packages, and now artificial Intelligence – can feel detached from a physical reality. However, a very real physical infrastructure is required for these services to exist.

This is made possible by the most complicated things humans make – the integrated circuit or the computer chip. Companies like Intel, TSMC, and Global Foundries make these chips, which are assembled into computers and storage devices that are literally the backbone of the modern economy. More recently, the Graphics Processing Unit, or GPU, which is specialized for rapid parallel multiplication operations, has been optimized for use in artificial intelligence applications in large language models.

These applications are expanding rapidly and, with easy access to the software and hardware needed for large language models, human culture is expected to undergo a disruptive transformation.

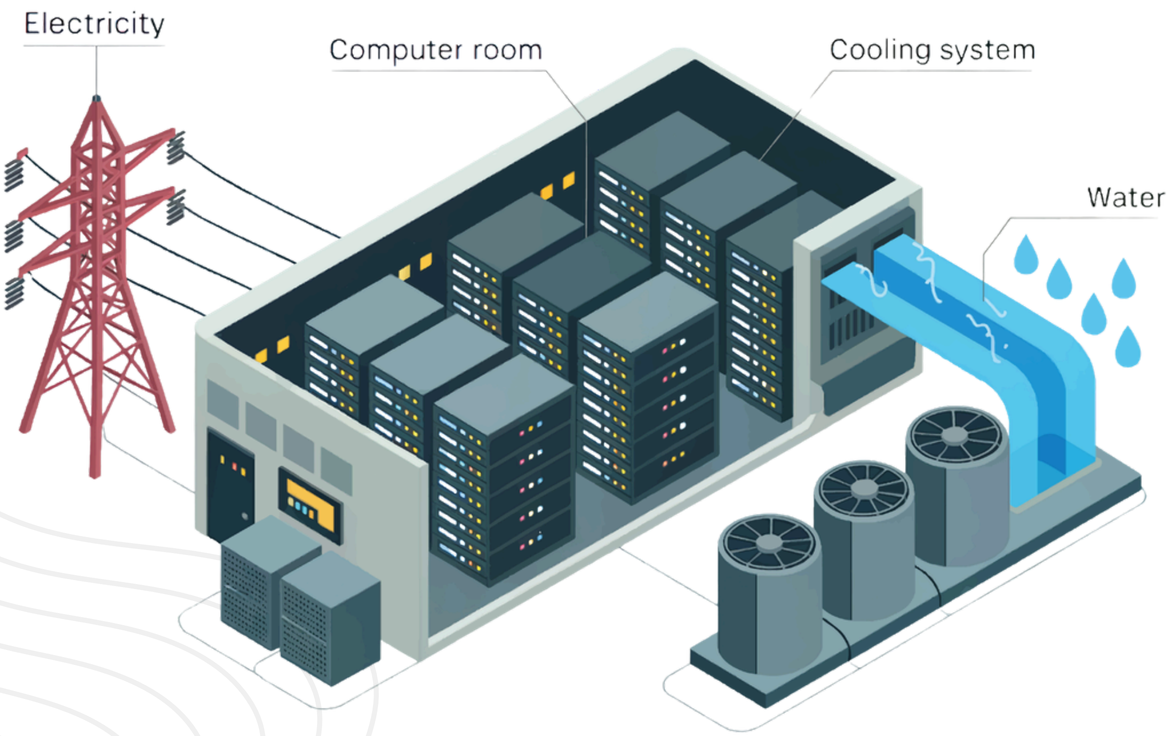

HOW DATA CENTERS WORK

Data centers, the places where computation occurs and where data is stored, require electricity to operate.

A measure of the rate of computational work being done in data centers is the rate of electrical work, or power, that goes into the data centers, which is measured in watts. The computational work to make a calculation or store and retrieve a bit of memory is actually quite small.

However, the number of calculations performed to generate responses is simply mind-boggling. Thus, while each calculation consumes only a little electricity and results in a tiny amount of CO2 released to the atmosphere, responses to searches or delivering streaming services or creating funny cat videos using Chat GPT or Grok require truly vast numbers of calculations.

The sum of the calculations consumes hundreds, even thousands, of megawatts in a data center.

To undertake these calculations, a voltage drop is applied across a circuit. This voltage drop induces a current to flow. When current flows, energy is dissipated and turned into heat. In other words, the circuit warms up. This resistive heating is the same as that which produces light in the old incandescent light technology.

Similar to incandescent bulbs, if circuits are not cooled, they will burn out. Thus, the megawatts required to generate calculations in a data center create heat, and the circuits must be cooled down.

The inputs to a data center are fiber optic cables that carry signals directing circuits to make the desired calculations, and the electrical power lines that are required for the computers to undertake the calculations. The outputs are a product – responses to the prompt – which exit through another set of fiber optic cables and heat equal to the energy input.

The computers are cooled in much the same way a motorcycle engine or an internal combustion car engine is cooled. For many motorcycles, air blows over the pistons, cooling them. In an internal combustion engine, water flows through the engine block next to the pistons, and it is heated.

That heated water is cooled by passing through the interior of a radiator at the front of the car where air is blown past the exterior of the radiator. The cooled water is recycled back to the engine.

Data centers are a bit more complex than a car engine, and therefore several cooling technologies are used. Each is based on lowering the temperature of the GPU by raising the temperature of a gas or liquid. That gas or liquid is subsequently cooled by external air. In this process, energy in the form of heat is transferred from the GPU to the outside air.

One process used in small data centers (less than 10 MW) is blowing cold air over the computer circuits. For the data centers used in artificial intelligence, these circuits are called graphics processing units, or GPUs.

Air cooling GPUs is very much like the air conditioning in a home. Cold air is blown across the GPUs, warming the air. This warm air is then blown across cold coils and is cooled down and returned to lower the temperature of the GPUs. In this process, the fluid in the coils is heated to the point of boiling. In fact, it is in the act of evaporating that heat is pulled from the air.

In the air conditioner cycle, the resulting gas is compressed, and the gas condenses. That produces heat, so the resulting liquid increases in temperature.

The fluid is cooled outside the house by blowing external air over another set of coils. (If you stand outside next to the air conditioner in the summer, you will feel heated air coming from the compressor unit.)

The condensed fluid is then returned to the coil in contact with the indoor air, which again causes it to evaporate. The house air is cooled, and the fluid evaporates again, thus continuing the cycle. Overall, heat from inside the house is transferred to the atmosphere outside the house.

There are variations on this cooling scheme. They each depend on a working fluid to extract heat from the GPUs. The working fluid is warmed and, to keep the cycle going, this fluid is cooled through a combination of mechanical work (as in the compressor used in air conditioning) and latent heat due to evaporation.

As data centers have grown in the power they draw, the heat generated by the racks of computers inside them has increased. For data centers consuming up to 80MW, chilled air removes heat from racks of computers, where each rack consumes as much as 8kW of electrical power.

Starting around 2021, and corresponding to the advent of AI and the use of GPUs, data centers began demanding 100-1,000MW, with each rack requiring 15-40kW. Under these conditions, air is no longer effective at cooling the computers. The air simply must flow very rapidly over the chips.

The result is a reversion to cooling technologies used in the 1960s-1990s for mainframe computers, in which water or an inert dielectric fluid is used to cool the computers.

For an open-loop liquid cooling system, cold water is passed through plates on which the GPUs are mounted. As it undertakes its calculations, the GPU warms. This heat is extracted by the water in the plate, with the result that the GPU is kept within its operating temperature range.

The hot water is then pumped to a cooling tower where it is sprayed into the air. This allows some of the warm air to evaporate and, in the process, chill the water. This water is returned to the cold plates to chill the GPUs. In this process, a great deal of water is evaporated.

Finally, there is a closed-loop system in which GPUs are immersed in a working fluid – water or another fluid. Immersion increases the rate of heat transfer, thus effectively increasing the rate at which the GPU can make its calculations. The working fluid is cooled in a closed loop by a refrigeration cycle.

If a closed-loop system is used, no water is lost in evaporation. However, the refrigeration cycle requires a lot of electrical energy, which can be more expensive for the operator than the cost of water. Often, the trade-off made by large data centers is between the consumption of water and the use of electricity.

Since 2022, with the introduction of ChatGPT, the demand for large-scale computers to undertake calculations leading to responses to prompts has driven up the demand for larger and larger data centers. The calculations are enabled by the large numbers of linked GPUs, and as the number of GPUs scales up, so does the electrical power going into the data center.

Data centers used for cloud computing are typically 50MW or smaller. The new data centers – often referred to as hyperscale data centers, are 100MW or larger.

UNDERSTANDING SCALE

How Big Is A 100-Megawatt Data Center?

1 million square feet ≈ 22 acres

The power use equivalent of cities like Boca Raton, Florida; Albany, New York; and Santa Barbara, California, which have populations of 90,000 - 100,000.

Relative to the size of a city, large amounts of electrical energy go into a data center which sits on a modest sized piece of land.

100MW results in 335,000 tons of CO2 released to the atmosphere each year, equivalent to driving 80,000 gasoline-powered cars for a year.

Energy Used For Data Centers In The United States

30 gigawatts of electrical power, the equivalent of 30,000 megawatts, is used to power data centers in the U.S., equivalent of the entire energy consumption of Australia, Spain or Taiwan.

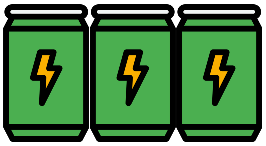

Electrical power demanded by data centers is expected to grow to 100GW by 2030.

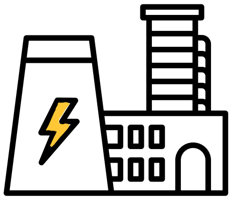

130 gigawatts represents a 12% growth in the total electrical power generation and use capability in the U.S.

The expansion is on the order of the peak power generation capacity of the United Kingdom.

While the power consumed by the data center can be very concentrated, the power itself may be drawn from geographically distributed sources. To ensure a stable source of power, hyperscalers prefer to draw from large grids that provide electricity for similarly large populations. Data centers therefore, become major users of the grid and thus major contributors to the carbon footprint of that grid.

THE ECONOMICS OF AI AND DATA CENTERS

There is a head-over-heels rush to build out data centers and the electrical and water infrastructure that will power them so that the U.S. can be in the lead of what is imagined as an AI enabled economy. Further, the federal government under Biden and Trump administrations see U.S. dominance in AI as a matter of national security.

In turn, communities are being approached to host 100-300MW data centers to meet this demand. In the rush, data center investments of $5-10B are not uncommon. Companies dominating this market (Google, Microsoft, OpenAI, Meta, Amazon Web Services) see enormous potential for profit in the application of AI.

In 2025, these hyperscalers spent approximately $350B on data centers and the projected 2026 spend is about $670B.

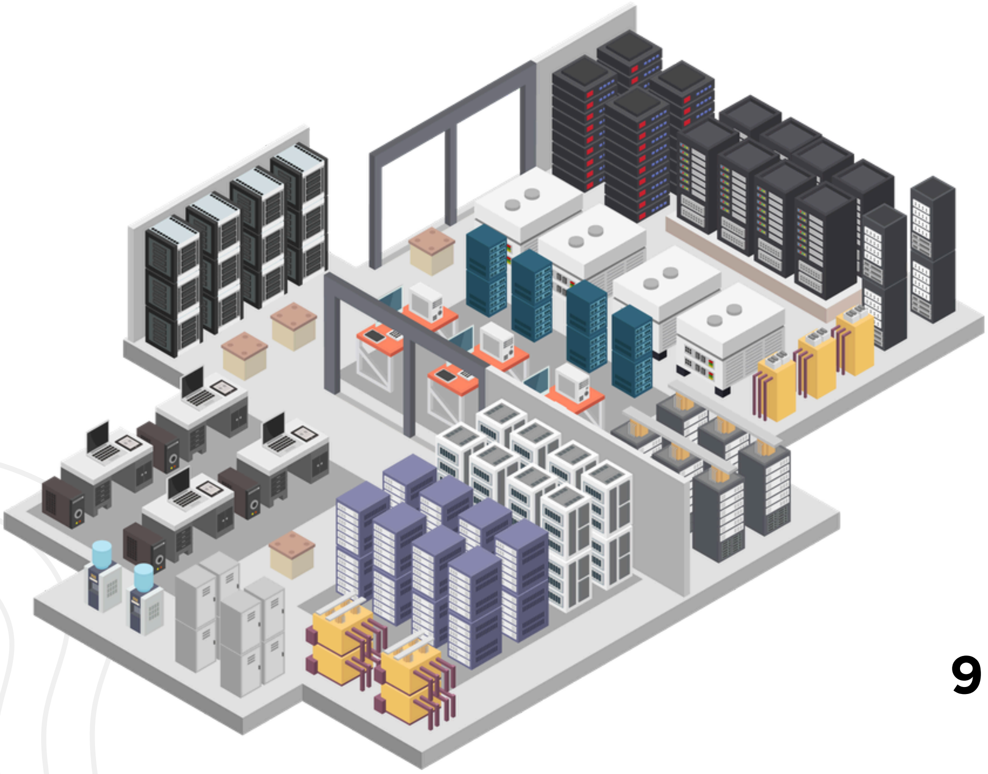

Development costs for a 100 megawatt data center

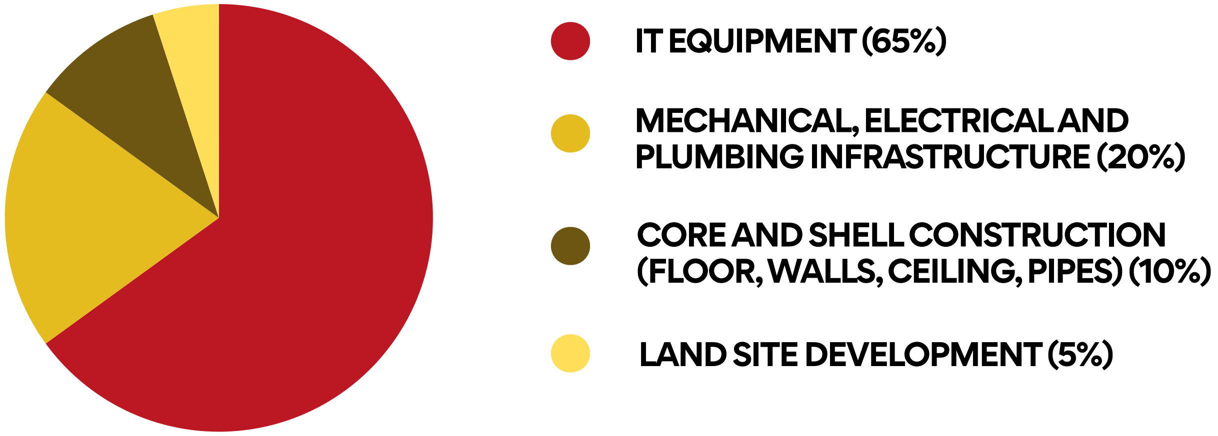

Total development expenditure at the upper end of cost ($~20M/MW): **$4.0B**

Approximate operating costs and profits for a 100 megawatt data center

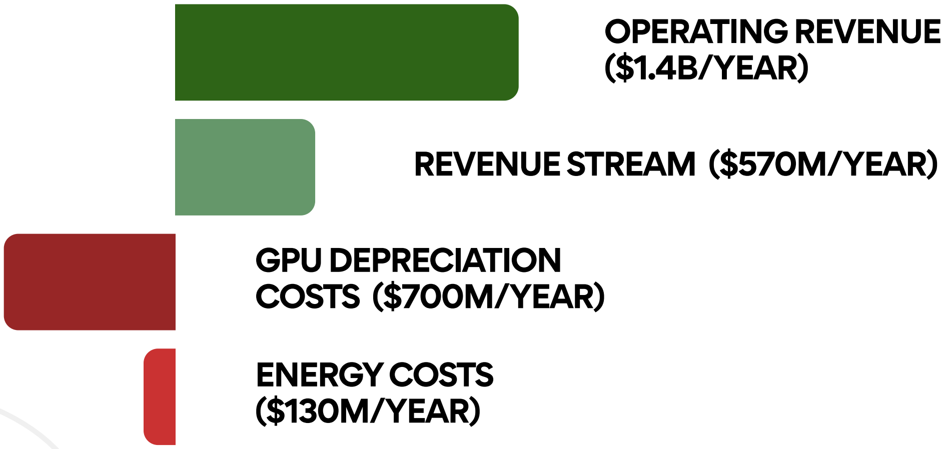

Hyperscalers are already the most profitable companies in the world because their AI data centers generate the computation that feeds every one of their high-margin businesses. Google converts computation into ads and AI Application Programming Interfaces, Microsoft converts computation into Office + Azure revenue, and Amazon converts computation into Amazon Web Services (AWS) and advertising revenue. In 2025, these three companies earned:

- **Alphabet: $132B**
- **Microsoft: $102B**
- **Amazon: $78B**

No other industrial sector on Earth produces profits at this scale. That is why there is extreme pressure to build more data centers — computational capacity directly drives corporate profitability.

Understanding these financial realities explains why developers often obscure the backers of data center developments. Market awareness may lead communities to ask for unwanted concessions from the deep-pocketed companies.

These profit margins also indicate why, in the face of growing protests at residential electricity rate increases tied to data center connectivity to the grid, in January of 2026, Google, Amazon, Microsoft and Meta signed agreements that would hold residential customers harmless for the cost of infrastructure required to connect their data centers with the grid.

Further, the industry has shifted from incentive seeking to a community-first posture. Microsoft has stated it will no longer accept tax abatements.

KEY QUESTIONS

Community members are often unprepared for the announcement of the development of a data center. The result is growing community pushback. In our research for the USC podcast, *Electric Futures*, which focuses on data center development in season 3, a series of common questions emerged that, if answered, can provide communities with information needed to make informed decisions. These queries involve rezoning, utility bills, environmental impact, including water use, noise levels, jobs, and economic benefits.

Our research and experiences indicate that each data center and each community are different. Thus, accepting at face value the hype of the data center developers or those opposed does not do service to individual projects. The details matter.

PROFITABILITY

Ultimately, these questions recognize that data center development results in large profits for companies that sell the results of computations. Communities should confirm that the developer has the financial backing and access to the market that will enable the data center to be profitable, and understand the consequences should the data center not be profitable.

Key questions on profitability:

- What is the business model of the data center? Will it be solely owned by a single hyperscaler? Will it lease its GPUs to the market in a co-location model?

- What are the current best practices of hyperscalers for tax abatements and holding residential customers harmless for electrical and water infrastructure costs associated with building and operating the data center? What is the developer asking for in this case?

- What is the expected return on investment for the investors in the data center and what are the sensitivities? (e.g., changes to the depreciation rate of GPUs can drive a swing from profit to loss)

- The developer and the community are assuming a great deal of the risk by allowing the project to go forward. How is that assumption of risk accounted for in community benefits?

- What is the expected utilization rate of the GPUs, and does the developer have upfront GPU time purchase agreements ensuring revenue streams to pay for building and operating the data center? How does the developer mitigate any risk that GPUs will not be fully utilized?

- How is the data center being financed? If the data center is speculative and being built on debt, what are the conditions placed on the developer by lenders, and why are financiers comfortable making that investment?

- What are the design constraints that impact the rate of return? (e.g., on annual operating expenses, what is the cost of depreciation and operation of a liquid cooling system and the cost of power as compared to GPU depreciation? What are the consequences of changes in market demand for electricity, water, and GPU-h? What are market trends?)

- The community is being asked to invest land, power and water resources, and those investments carry risks that the market may change (i.e., is there an AI bubble?) How is that investment and willingness to take on risk recognized in community benefit agreements sought by the community with the developer?

COMMUNITY ECONOMIC BENEFITS

Data centers building involves a substantial number of skilled construction workers.

As a rule of thumb, 5-8 skilled construction jobs are required per megawatt in a construction period that is typically 18-36 months. For a 100MW facility, there will be 500-800 full-time equivalent workers with peak jobs during the fit-out (when mechanical and electrical systems are installed; ~1000 workers).

Estimates suggest that for each construction job at a data center, there are 2.5-3 indirect fobs in the local supply chain – for example, in the concrete batch plants, local transportation and service sectors. There is currently a shortage of skilled labor needed to meet the explosive demand for data centers such that those with the desired skills are seeing increased wages and steady work – if workers are willing to move to where data centers are being built.

Once complete, there is a substantial drop in labor demand. Data centers are essentially “dark factories” where manufacturing of value-added products is done without substantial human intervention. For each megawatt, there are 0.5-1 permanent jobs produced. Thus, a 100 MW facility will have approximately 50-100 full-time employees.

There is growing evidence that the hosting county or city does not see general increases in per capita GDP or per capita income. This is attributed to the limited number of jobs associated with the operation of the data center.

Data centers are subject to taxes and fees used to support the local community. Their magnitude, and the exemptions given as incentives to attract businesses, are local considerations involving state, city and county taxing ability.

Increases in tax base and revenue are often cited as a major advantage to hosting a data center. The magnitude of the revenues generated by these taxes depends heavily on the incentives afforded to the data center to bring it to the community.

A community considering hosting a data center should explore both the actual taxes and the incentives (for example, a common incentive for the community is to forego property taxes for 20 years).

Tax categories

- **Sales, Use, and Transaction Privilege Tax** (mostly in Arizona – construction and equipment)

Identify what is taxable vs. exempt in the hosting region; estimate local share if applicable. (That is, who collects the tax – city, county, state – and how the community benefits from that payment or lack thereof.)

Sales tax is paid at the point of delivery of the product. Thus, the builder of the data center would pay, for example, a sales tax on a GPU that constitutes 60% of the cost of building out the data center. This is a one-time tax, but it drives up the cost of ownership. Many states and communities will forego sales tax as an incentive to attract data centers.

- **Property Tax** (real and personal)

Taxing authorities should identify the expected new assessed value, note any abatements and fees, and benchmark against high-yield jurisdictions. Real tax is on the value of land, while personal tax is on the property in the data center. By abating the personal property tax for the data center for, say, 20 years, the value of the equipment in most facilities has been greatly depreciated so that once the tax is applied, the equipment is valued at a lesser amount than when it is new.

Data centers are unique in that the lifetime for the GPUs is short (3-6 years). Thus, the replacement rate is high, so that there is limited depreciation of the value of this property. Communities can forego large sums by not collecting property taxes on the data center and the computer equipment inside. These taxes are paid by the data center owner.

- **Utility and Access**

Confirm franchise fee rates exist, learn how they are distributed and abated, and determine applicable system development charges. One-time fees may help cover capital costs of expanding utility systems, typically applied to water, sewer, storm, and transportation systems. These are a source of revenue for communities to sustain their infrastructure.

Data centers can rapidly dominate the use of this civil infrastructure. Ensuring the load is reasonably distributed will be important to the community.

<u>Incentives</u>

- States and communities compete to attract businesses. Often, incentives are given to locate in a particular geography. In some localities, these incentives are directed explicitly at data centers while in others, they are general-use vehicles such as enterprise zones. Enterprise zones are areas established typically within counties or cities with agreement with the state government, where taxes can be abated in efforts to attract business. These tend to be very local, with state, county and city capacity to establish incentives.

- Large enterprises, including data centers, often accept incentives and then reach an additional agreement with direct payments in lieu of taxes and/or to support communities. These can be one-time or recurring direct payments to the communities and are codified in Community Benefit Agreements (which can be legally binding), explicitly supporting schools and infrastructure (such as water handling infrastructure).

As an example, the data center may pay a fraction (25%) of foregone property taxes. Community Benefit Agreements can be negotiated with community groups, for example, NGOs and unions, but can also be negotiated with the local governments to ensure that specific community needs are met as the data center is built and operated.

Key questions on community economic benefits

- How many jobs will be created in the building of the data center and what is the duration of those jobs?
- What are local sourcing agreements for construction and prevailing wage and trade contracts?
- How many jobs will be available once the development of a data center is complete?
- What incentives are being granted to the data center by the state, county, and city?
- How much revenue will be returned to the community, and how will it be used?
- How many permanent jobs will be created by the data center and what is the number of indirect jobs expected? What is the nature of these jobs and the backgrounds required?

ENERGY USE AND UTILITY BILLS

Data centers use large amounts of energy to store and distribute information, to manipulate information, and to perform the mathematical calculations that are the basis of AI. They require a continuous supply of electricity and pay the costs for multiple levels of redundancy.

Typically, hyperscalers seek to connect data centers to an existing grid to provide access to a large and well-regulated electricity distribution network. The utility is the distribution network, but it often does not own generation capabilities. Utilities must furnish “adequate, efficient, and reasonable service,” and cannot pick and choose only profitable customers. If a customer seeks more power, the utility has a responsibility to supply it. This is done by holding auctions where the utility seeks suppliers of energy who bid on the cost at which they will supply the energy. The utility signs a purchase power agreement with a supplier that is often decades long.

The utility establishes a rate structure at which it sells this electricity. Often, the utility is a regulated monopoly, and a commission sets the profit of the utility and approves the pricing structure with profits on the order of 10-15%. The utility passes on the cost of the electricity directly to the customer, such that if gas prices fluctuate or there is a spike in gas prices due to a winter storm, the resulting costs are transferred to the customer.

The profit of the utility is calculated on the basis of their capital equipment – transmission and distribution networks, substations, transformers – which they are incentivized to grow as their profit depends on the depreciating capital assets.

Electricity infrastructure is fundamental to the smooth operation of society and the economy. As a result, the power grid is heavily regulated. The infrastructure is expensive and is amortized over decades. Plans thus look far into the future.

Over the past two decades, electricity consumption in the US has grown only slightly. As a result, the industry has been able to spend modestly on infrastructure development, mostly replacing infrastructure that has reached the end of its working life.

The sudden expansion of data centers with their enormous electricity demands places the utilities in a bind. The utilities are obligated to connect to the data centers and they are incentivized to do so as the capital expenditures associated with the resulting electricity demand will increase profits. At the same time, the demand is occurring faster than the market can respond to generate the needed electricity and to build the necessary transmission and distribution backbone.

Utilities use a price-setting model that distributes the costs of new infrastructure across all users. This enables new developments to come online without taking on the full cost of the infrastructure the new development demands. As data centers demand for electricity expands and large amounts of infrastructure are built, residential and small businesses are hit by substantial increases in electricity rates, resulting in a growing backlash.

To increase local acceptance, data centers are being encouraged to bring their own electricity. Essentially, they become their own power generators, which increases permitting requirements, but also allows data center developers to move faster than the utilities are able. Increasingly, utilities look at the backup and operational capacity of the data center as a resource to use during an emergency such as a heatwave or a winter storm.

An agreement between the utility and a data center can include a certain number of hours that the utility can tell the data center to reduce its demand on the grid, thus releasing power to be used in other places. The data center responds by performing less computation, and/or it can resort to backup generators to replace reliance on the grid. Backup generators are kept on site to ensure redundancy in access to power. They are most often diesel engines with substantial capacity to generate soot – an environmental hazard – and they are able to generate two times the average power use of the facility.

Key questions about power

- How much power will the data center use? How does this compare to the power usage in the community? How does this power demand compare with the average and peak power demands of the utility that will supply the power?

- How will the connection of the data center impact the utility's reserve power? When power consumption is curtailed during peak demand times, will the data center cut back, or will the small businesses and residential customers be required to cut back?

- How will connecting the data center to the local grid impact the stability of the grid and the size of the reserves that are available during heat waves or winter storms?

- What infrastructure build-out is necessary to connect the data center to the local grid? Who pays?

- If a hyperscaler states that it will pay its fair share of grid upgrades, how is "fair share" defined?

- Are residential and other commercial energy bills anticipated to rise because of energy use by the proposed data center?

- Are power purchase agreements with the utility in place, and how will those impact the cost of residential and small business electricity?

- Will the data center build its own electricity production capability and operate behind the meter? What is the source of that electricity? How will it interface with the grid?

- What are the agreements with the utility about how often the backup generators on site can be used to generate electricity that the data center will sell to the utility, and how are those prices set?

- What are the state mandates for carbon-free electricity? How will the construction of the data center impact the ability of the state to meet these obligations? If the data center is not developed by a hyperscaler, what are the commitments of the developer for carbon-free electricity usage and are those contractual obligations?

- Does the agreement with the data center and the utility require that the data center's usage of electricity be reported?

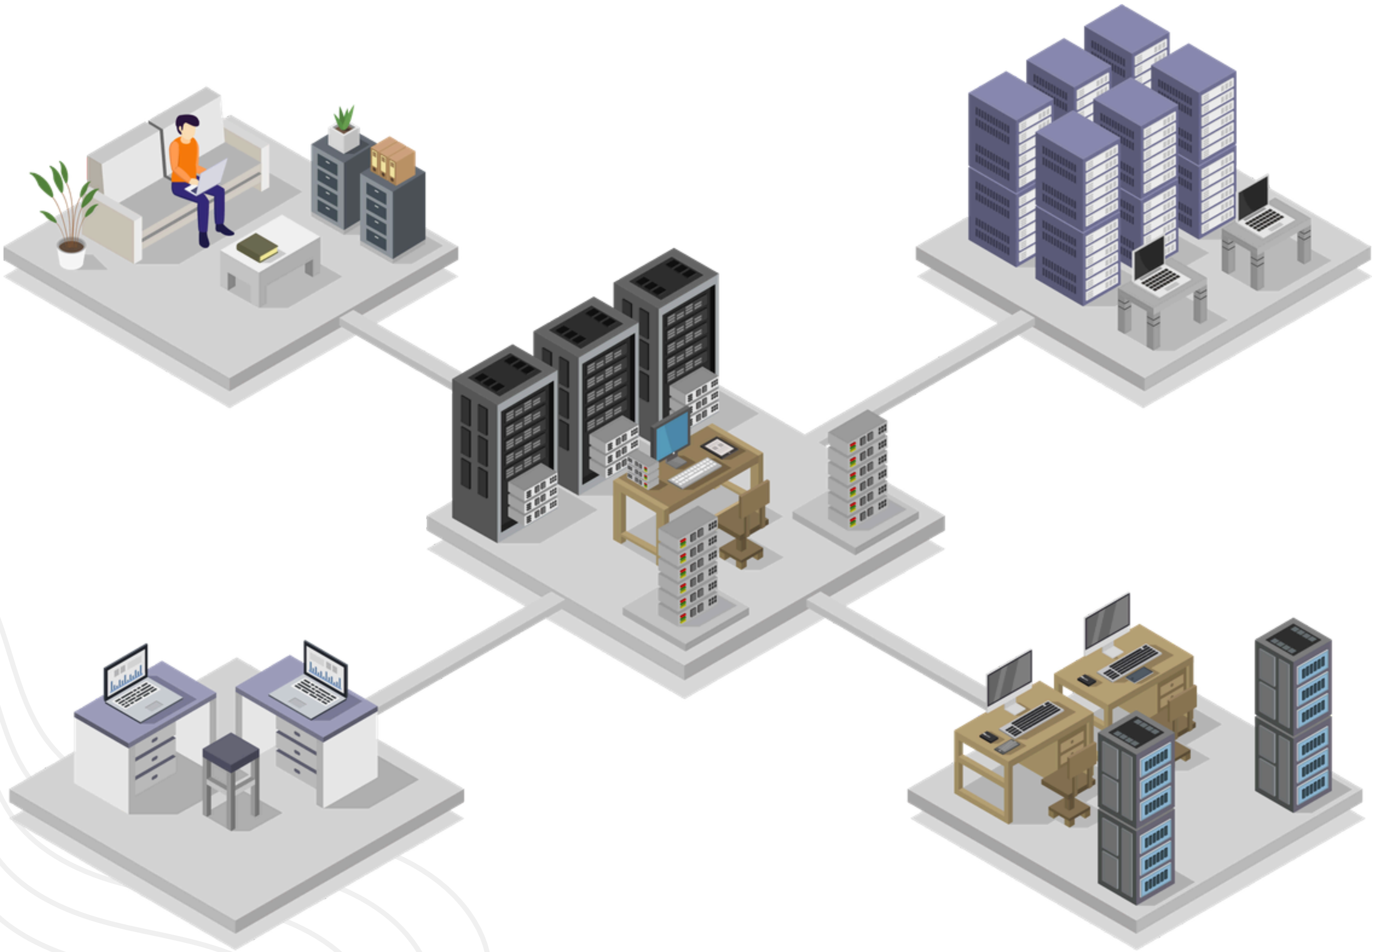

WATER

Data centers take in megawatt hours of electrical energy and convert this energy to heat. If the data center is not cooled, it will cease functioning. Early data centers used air cooling. Higher-density data centers (50-150MW) use water cooling. Larger data centers (300MW-1.5GW) require direct-to-chip, or immersive cooling.

Water is used for cooling because it is cheap and it has a large latent heat of evaporation. In other words, water provides a lot of cooling for a little evaporation. The figure of merit for water usage is the Water Use Effectiveness, WUE. This is the liters of water used per kWh of electricity used in computation. WUEs of waste water cooled data centers are 1.8 or greater. A 100MW data center with a WUE of 1.8 will consume about 400 million gallons of water a year, or about 1300 acre feet.

Aggregated water use data is reported by hyperscalers in their annual reports. However, accessing water usage data at a particular facility is often difficult, as that information is held as proprietary. Following a lawsuit, Google released 2024 water usage information for its data centers in The Dalles, Oregon, where the centers use 150-200MW of power with a WUE of 1.9-2.2 and consume 434M gallons of water, ~30% of the local utility's potable water production.

Many data centers use potable water from the community's water utility. If the city hosting the data center is small, the data center can consume 30-50% of the utility's total capacity. In times of low rainfall or drought, this places a strain on the data center and the community as to who has primacy of rights for that water.

Water usage for cooling has been implicated in a number of aquifer issues associated with the concentration effects of evaporative cooling.

In Umatilla County, Oregon, industrial farming has resulted in increasing nitrate concentrations in the aquifer for the past several decades. When aquifer water is used for data center cooling, nitrates are concentrated such that returning the unused water to the aquifer accelerates the increase in concentration in aquifer water, resulting in serious health concerns.

Total water usage by a data center is not limited to the water used on site. The largest secondary use of water comes from power generation. If electricity is generated by burning fossil fuels, water is often used to cool the power plant, creating the electricity that flows to the data center. This can increase water usage attributed to the data center by 50%.

Key questions on water use

- What is the cooling technology that will be used in the data center?
- What is the projected Power Use Effectiveness of the data center and how does it compare to industry best practices?
- What is the Water Use Effectiveness of the data center and how does this compare to industry best practice?
- If water is used for cooling, what is the source of that water?
- How much water will be used by the data center and how does this compare with the amount of potable water currently being delivered to the community?
- Do agreements with the hyperscaler require the data center to report water usage?
- What is in the water that will be concentrated in the cooling process and will enter the aquifer in concentrated form after the water is used for cooling?

NOISE

Data centers are industrial sites. Massive amounts of energy are consumed, and massive amounts of heat are expelled. Each of these operations involves fans, pumps, transformers, and backup generators. These operations are all possible sources of noise.

Noise propagates large distances and is perceived differently depending on the time of the day, atmospheric conditions (humidity, temperature, atmospheric pressure), and frequency range. Noise level ordinances are controlled at the local (city and county) level. Thus, within the limits of fair treatment, the local authority can demand that certain levels are met.

Surprisingly few data centers have seen controversy over noise levels, and those where noise has been an issue have not persisted. That is, the issues appear to have been resolved.

Testing of backup generators is a sporadic noise event and can be problematic, but perhaps the larger issue with these generators is the soot levels that are generated during their operation.

Data centers can be designed to mitigate noise concerns. These include enclosure of noise-making equipment, walling the property, and creating setbacks from the property line with landscaping that muffles and dissipates the noise. Noise studies can be carried out to confirm that the design will mitigate noise transmission into surrounding neighborhoods.

One study attempting to determine if living close to a data center was sufficiently unpleasant to lower property values found that in Loudon County, the location with the highest density of data centers in the country, there was a weak correlation that showed, if anything, a slight positive trend with property values increasing closer to a data center.

This study was partially motivated by complaints about data center noise in Loudon County and a desire to find evidence of community responses. The result suggests that the sound issues can either be mitigated or that some people are more sensitive to the noise than others.

Key questions about noise

- What are the local noise ordinances?
- Do local ordinances reflect the best practices for sound mitigation from data centers (for example, decibel levels in particular frequency ranges)?
- What is the design of the landscaping around the data center to dampen noise propagation (offsets, enclosed diesel generator farms, fans at ground level, vegetation and berms)?
- Are sound studies required, and is regular monitoring part of permitting?
- Are there consequences if the data center exceeds local ordinances?

SUGGESTED “BEST-PRACTICE” TAKEAWAYS FOR HOST COMMUNITIES

These recommendations are based on takeaways from case studies that are highlighted in Appendix B.

- Codify siting & performance early (Marana): define “data center,” require rezoning/Specific Plans, and bake in water source constraints, power availability proofs, and noise studies at entitlement.

- Treat backup-diesel as a cumulative risk problem (Quincy): require publicly noticed air permits with health-risk modeling, stack design, testing limits, and community hearing cycles.

- Use formula-driven “in-lieu” payments (Oregon SIP/CSF): predictable 25% of abated tax (cap) paid annually; allow local add-ons directed to water/energy infrastructure and rate mitigation.

- Bind operational limits through enforceable agreements (Saline Township): consent/development agreements (or CBAs) can lock in no-evaporative-cooling, noise caps, setbacks/berms, open-space easements, and community funds.

- Update noise regimes to capture low-frequency/HVAC + generator testing (Prince William/Loudoun): modernize measurement (C-weighting or octave bands), limit testing hours, and require acoustic modeling at review.

APPENDIX

A. WHY IS THERE A RUSH TO BUILD DATA CENTERS NOW?

The short answer to this question is the penetration of AI into the economy. The even shorter answer is the profit derived from the data centers.

Below is an analysis of the costs to build and operate a 100MW data center. This provides a sense of the profitability of the data center and the risks taken by the data center developer as they create their business case. The bottom line is that the data centers are enormously profitable lines of business.

In addition, they are the underlying infrastructure for the large software suppliers (Google, Meta, Oracle, Amazon). These hyperscalers purchase computational capabilities from the data center and then sell the results at substantial markups to the ultimate consumer. The companies are betting AI will open vast new markets and generate huge profits, thus driving a rush to gain access to the infrastructure that opens the door to highly profitable revenue streams.

Build Costs / Capex per MW
- Turner & Townsend — Data Centre Cost Index (2024)
- Turner & Townsend — Methodology (what's included/excluded)
- Cushman & Wakefield — U.S. Data Center Development Cost Guide (2025)
- McKinsey — Building data centers bigger, faster, cheaper (Aug 2025)

Power & Energy Cost Benchmarks
- EIA — Electric Power Monthly (Table 5.6.A, industrial $/kWh by state)

AI Hardware Power / Rack Density
- AI server draw & rack density (5–7 kW per 8×H100)

GPU Pricing & Cloud Rental Rates
- H100 rental price comparison (multi-provider)

Facility-level Profitability / Economics
- Thunder Said Energy — Data-centers: the economics (capex/opex model)

Corporate/Cloud Profitability & Growth Context
- CRN — Q3 2025 Cloud earnings face-off (AWS, Microsoft, Google Cloud)
- Visual Capitalist — Most Profitable Companies (2025 ranking)
- IEEE ComSoc Tech Blog — 2026 hyperscaler capex > $600B

Depreciation Lives & Lifecycle (book vs economic)
- Bytelota — GPU depreciation crisis / 6-year lives debate
- SiliconANGLE — Resetting GPU depreciation (reuse: training → inference)

B. EXAMPLES OF COMMUNITY ORDINANCES AND AGREEMENTS RELATED TO DATA CENTERS

Codifying standards

Marana, AZ

Marana is a small community to the west of Tucson. Anticipating that there might be interest in bringing a data center to Marana, the Mayor undertook a study to understand controls that might result in ease of negotiations and meet citizen needs. The result is what appears to be best practice standards for power, water and noise.

A 550-750 MW data center has now been approved to occupy a 661-acre parcel of land.

Key elements include:

- No potable water for cooling (developers must secure a non-potable/alternative water source for cooling/humidity control)
- Utility proof-of-power required
- Acoustical study and noise standards
- Public rezoning hearings to site a principal-use data center

Developers are required to announce extensive information ahead of development, and final occupancy permits are dependent on meeting acoustic standards. The developer must have commitments from utilities that they have enough water and power to supply the facility, and will have to report energy needs projections and annual water consumption.

Back-up Diesel Generator Management

Grant County, WA

Grant County in Washington State is home to multiple data centers with a combined electricity draw of 500-700MW.

Diesel soot produced during the periodic testing of backup diesel generators is a considerable point of contention. Schools are approximately a mile from the diesel generator farms, but some of the data center campuses are surrounded by residential developments. The result was a great deal of study and mitigation of the testing of these diesel generators.

A health risk analysis was conducted by the Washington Dept of Ecology, and a discussion of mitigation undertaken (includes when diesels are turned on, how many hours a month they can be operated and how many can be tested at any given time).

Best Practices in Enterprise Zone and Strategic Investment Program, SIP, and Community Service Fees.

Wasco County/The Dalles, OR

Google built its first self-owned and operated data center in The Dalles in 2002. Since that time, the data center capacity has grown to 200- 350MW. Issues were water infrastructure and abatement of property taxes. The result is a series of agreements with regard to water and support in lieu of taxes.

Controversy over expanding data centers and water use is a current topic of discussion in the community.

Court Enforceable Community Driven Agreements

Saline Township, MI

Oracle, OpenAI and Related Digital plan a 1.4GW data center in Saline Township. Considerable public backlash resulted in a court-binding agreement that stands essentially as a Community Benefits Agreement. Included in the agreement are:

- Mandatory buffer zones and preservation of wetlands
- Noise limits (facility noise cap; testing restrictions)
- Closed-loop cooling
- A battery energy storage system (~$2B)
- A $14M fund for community benefit

The large energy requirements of this data center have resulted in concerns by environmental groups, who have seen no commitments to ensure Michigan's clean energy goals will not be postponed by this project.

Noise Ordinances

Prince William County, VA

Updated noise ordinance after a year-long technical review driven by resident complaints about low-frequency HVAC hum & generator noise near neighborhoods. The county also commissioned a noise-consultant scope specifically calling out data center construction/operational/generator noise and mitigation methods. There is disagreement as to whether the noise standards meet community needs. Noise levels continue to be a matter of discussion.

Rules developed after years of Data center operation

Chandler, AZ

In addition to the above examples from Virginia, where their buildout of data centers results in continued changes to local ordinances, Chandler is a good example of how approaches have evolved.

Chandler welcomed data centers in the 2000’s – 2010’s. neighbors complained about the low-level hum generated by the data centers. In 2023, Chandler adopted a city ordinance that established conditions for hosting data centers.

This resulted in a data center capacity of 250-300 MW. More recently, Chandler has closed its doors on further data center expansion, seeking instead economic development that brings jobs to the community. Nevertheless, their ordinances provide a template for dealing with issues arising when hosting data centers.

Chandler’s approach has been adopted in nearby Phoenix.

C. ADDITIONAL BACKGROUND AND RESOURCES

Operating costs related to data centers

Operating costs include electricity at the price negotiated by the data center with the utility. A nationwide average is $ 0.085/kWh (but can be 3-4 times this at different locations. Staffing and maintenance are estimated to cost ~3% of the non-GPU capital expenditure. Using a typical 70% utilization rate of the GPUs (the fraction of time they are making computations) and accounting for 20% more power not going to the GPU's to operate the facility (with much of this going to cool the facility) yields an electricity bill of ~$63M and $60M for maintenance and staff costs. On top of this is the ~$700M that is amortized to take into account the need to replace the GPU's in three years and the building in 20.

Data centers make money selling time on their GPUs. This is typically priced at $3.75/ GPU-h. The result is that 60,000 GPUs operating at a 70% utilization rate yields an operating revenue of ~$1.4B/yr.

Thus, the data center is expected to bring in $1,400M/yr with an operational cost of $830M/yr, yielding a revenue stream of $570M/yr.

These calculations indicate that assuming the GPUs are fully used at the 70% utilization, a 100MW data center can throw off $570M/yr and still replace the GPU's every 3 years. A check on this calculator is that data center earnings after depreciation but before interest and taxes are typically reported as $4M-$10M/MW. The above calculation yields ~$6M/MW.

Over the next 5 years, the US expects to build out an additional 50GW of data center power consumption, suggesting a growth in earnings of $300-600B/yr. The 2025 power draw for data centers in the US is ~24GW.

Thus, the 50GW will take the nation from 24 to ~75GW, tripling the revenue. However, this is simply market growth. The calculations made here do not project profitability growth with the penetration of AI.

The rate at which a data center is paid for its GPUs is set by a competitive market. That market holds down the price required to buy computations. If the hyperscaler owns the data center, they set their own depreciation rate. In the last couple of years, this has been stretched from 3 to 6 years, essentially cutting in half the $600M/yr in GPU depreciation in the 100MW data center example, lowering operating cost plus depreciation from $700M/yr to $400M/yr an increasing operating margin from ~$600M/yr to $900M/yr. This dramatically reduces the payback time for a data center from years to quarters, making the data center even more profitable.

Secondly, the market price of a GPU-h reflects the cost of delivering the computation and market competitive forces. This cost does not reflect the price charged by the hyperscaler for the services they offer. The hyperscalers are expecting that the AI services they offer will be valued at 4-5x the current revenue per MW. In essence, while the hype scaler pays $6-$12/GPU-h, it sells the computation at closer to $30-$50/hr, thus generating a much larger revenue, again driving up the profitability of having access to a data center.

The price per GPU-h is also set by demand. Should there be a slowdown in demand for GPU computation, the revenue could drop dramatically as data centers seek to continue operation.

In addition, as we look into the future, the stranglehold Nvidia has on the GPU market is unlikely to be sustained, such that competition will drive down the cost of the GPUs.

Thus, while data centers are currently highly profitable, changes in the market could drive that profit down or up. Those betting tens of billions of dollars on data centers are taking substantial risks. So are the communities that host the data centers.

Depreciating costs at data centers

The physical assets are depreciated over 20 years, yielding an operating cost of $100M/yr, which must be accounted for. There is controversy over the time period during which the GPUs will have a value-generating capability. The standard in many discussions is that the GPUs are depreciated over 3 years - that is, the operations must pay for the GPUs in three years so that they can be replaced by the next generation.

The hyperscalers have pushed this to 6 years, thus cutting in half the revenue needed to pay down the GPUs and significantly increasing data center profitability - as long as they do not replace the GPUs in 3 years instead of 6. I will stick to 3 years depreciation of the GPUs. This means the data centre's operations must generate $600M/yr to account for GPU depreciation.

Community economic benefits

- How Constructing New Data Centers Will Impact the American Economy
- Data Centers Are a ‘Gold Rush’ for Construction Workers
- Data Center Employment Forecast Analysis
- Why Data Centers Will Be Economic Development Duds
- What Happens When Data Centers Come to Town?

Incentives

- Examining the impacts of Virginia’s data center industry and site location criteria
- Extracting Profits from the Public: How Utility Ratepayers Are Paying for Big Tech’s Power
- What Happens When Data Centers Come to Town?
- Tax Revenue & Economic Development Subgroup Findings
- Washington's hydropower has created a data center boom. Some are concerned about its future

Energy use and utility bills

- Decarbonizing the Obligation to Serve
- Electric Cost Allocation for a New Era
- 2024 United States Data Center Energy Usage Report
- As electric bills rise, evidence mounts that data centers share blame. States feel pressure to act.
- How families could get stuck with higher electric bills if the AI data center boom goes bust
- Behind-the meter generation is scaling up to meet “hyperscale” US demand

Water

- Uptime Institute Global Data Center Survey 2024
- Data Centers and Water Use
- PUE vs. WUE: Balancing Efficiency & Sustainability in Data Centers
- Calculating the Impact of Water Usage on Data Center Costs and Sustainability

Noise

- Data centers challenge communities: revising noise ordinances for balance
- Sound Solutions for Data Centers: A Guide to Noise Control
- Virginia's 'Data Center Alley' residents say an eerie hum is keeping them up at night
- Amazon Tones Down Its Data Center Noise After Residents Sound the Alarm
- Prince William Supervisors greenlight new noise ordinance

USC
Annenberg
Center for Climate Journalism
and Communication