CoreWeave

AI Wars 2025: Power, Infrastructure, and the New Race for Scale

Why Data Centers, Not Algorithms, Will Determine The Winners of the AI Wars

---

Skidmore College – Capstone Executive Presentation
Landon Schwartzman, Robyn Kachere-Sun, Stephen Dumaresq, Grant Steinberger
Fall 2025

CoreWeave

CoreWeave

01 Industry Overview

The AI industry in 2024 was valued at $638.23 billion and is expected to grow at a CAGR of 19.2% in the next 10 years

Notable Market Highlights:

- North America has the largest market share, 36.42% in 2024
- Asia Pacific region is expected to grow at a CAGR of 19.8% from 2025 to 2034
- Solution type, software segment has the biggest market share of 51.4% in 2024
- Solution type, services segment is projected to grow at 18.3% from 2025 to 2034
- Technology segment, machine learning has a market share of 36.7% in 2024
- Technology segment, generative AI is expected to grow 22.9% from 2025 to 2034
- Functions segment, cybersecurity is expected to grow at 20.4% CAGR from 2025 to 2034
- The healthcare segment is expected to grow at 19.1% over the projected period

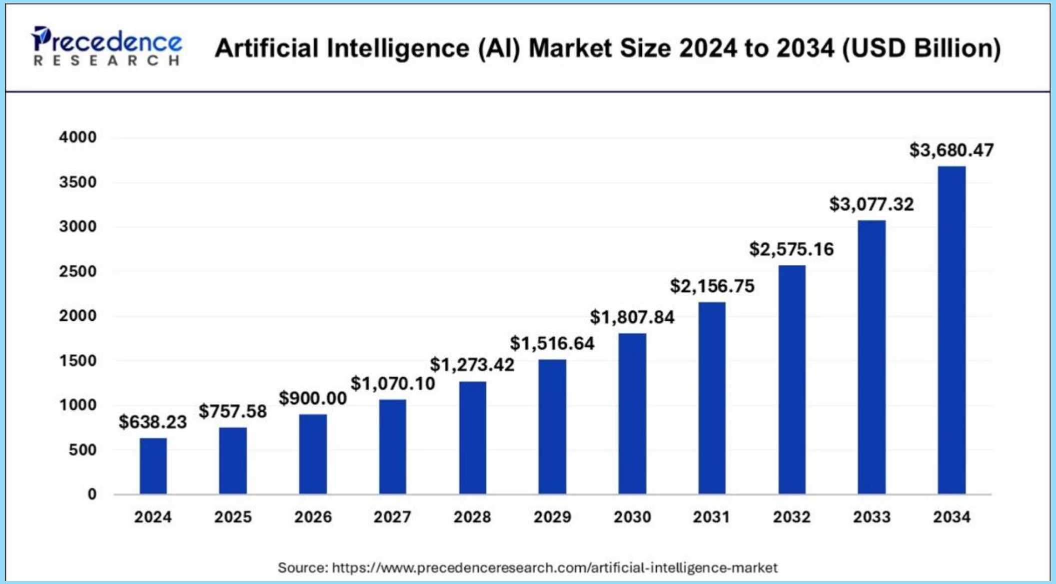

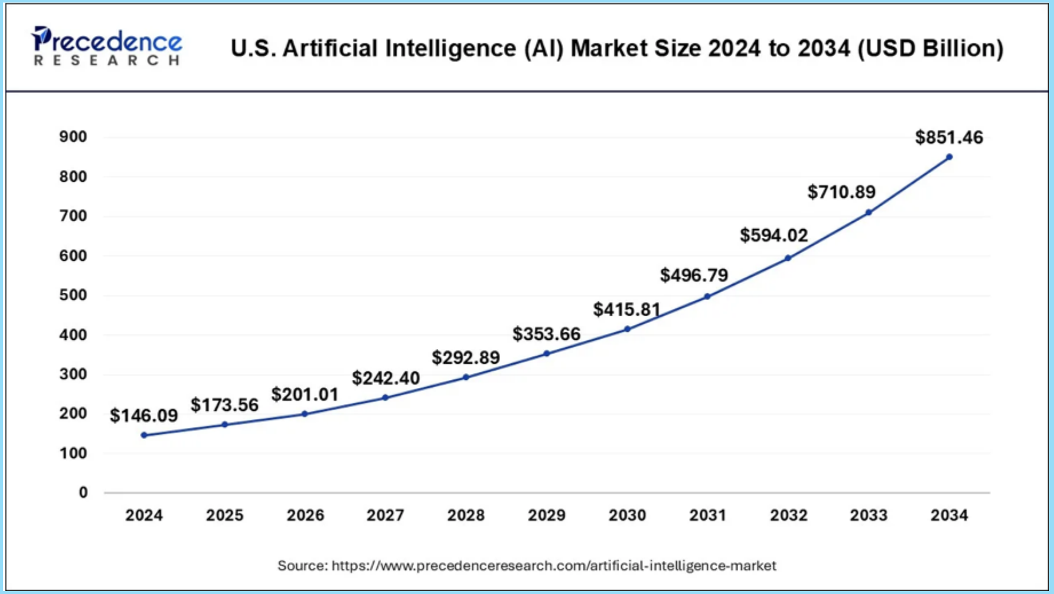

Key Takeaway: North America leads the market, software dominates, machine learning is the core technology. Generative AI, cybersecurity, and healthcare are the fastest growing segments.

Source: Precedence Research

There are several types of AI players –from chip manufacturers, and cloud providers, to AI companies

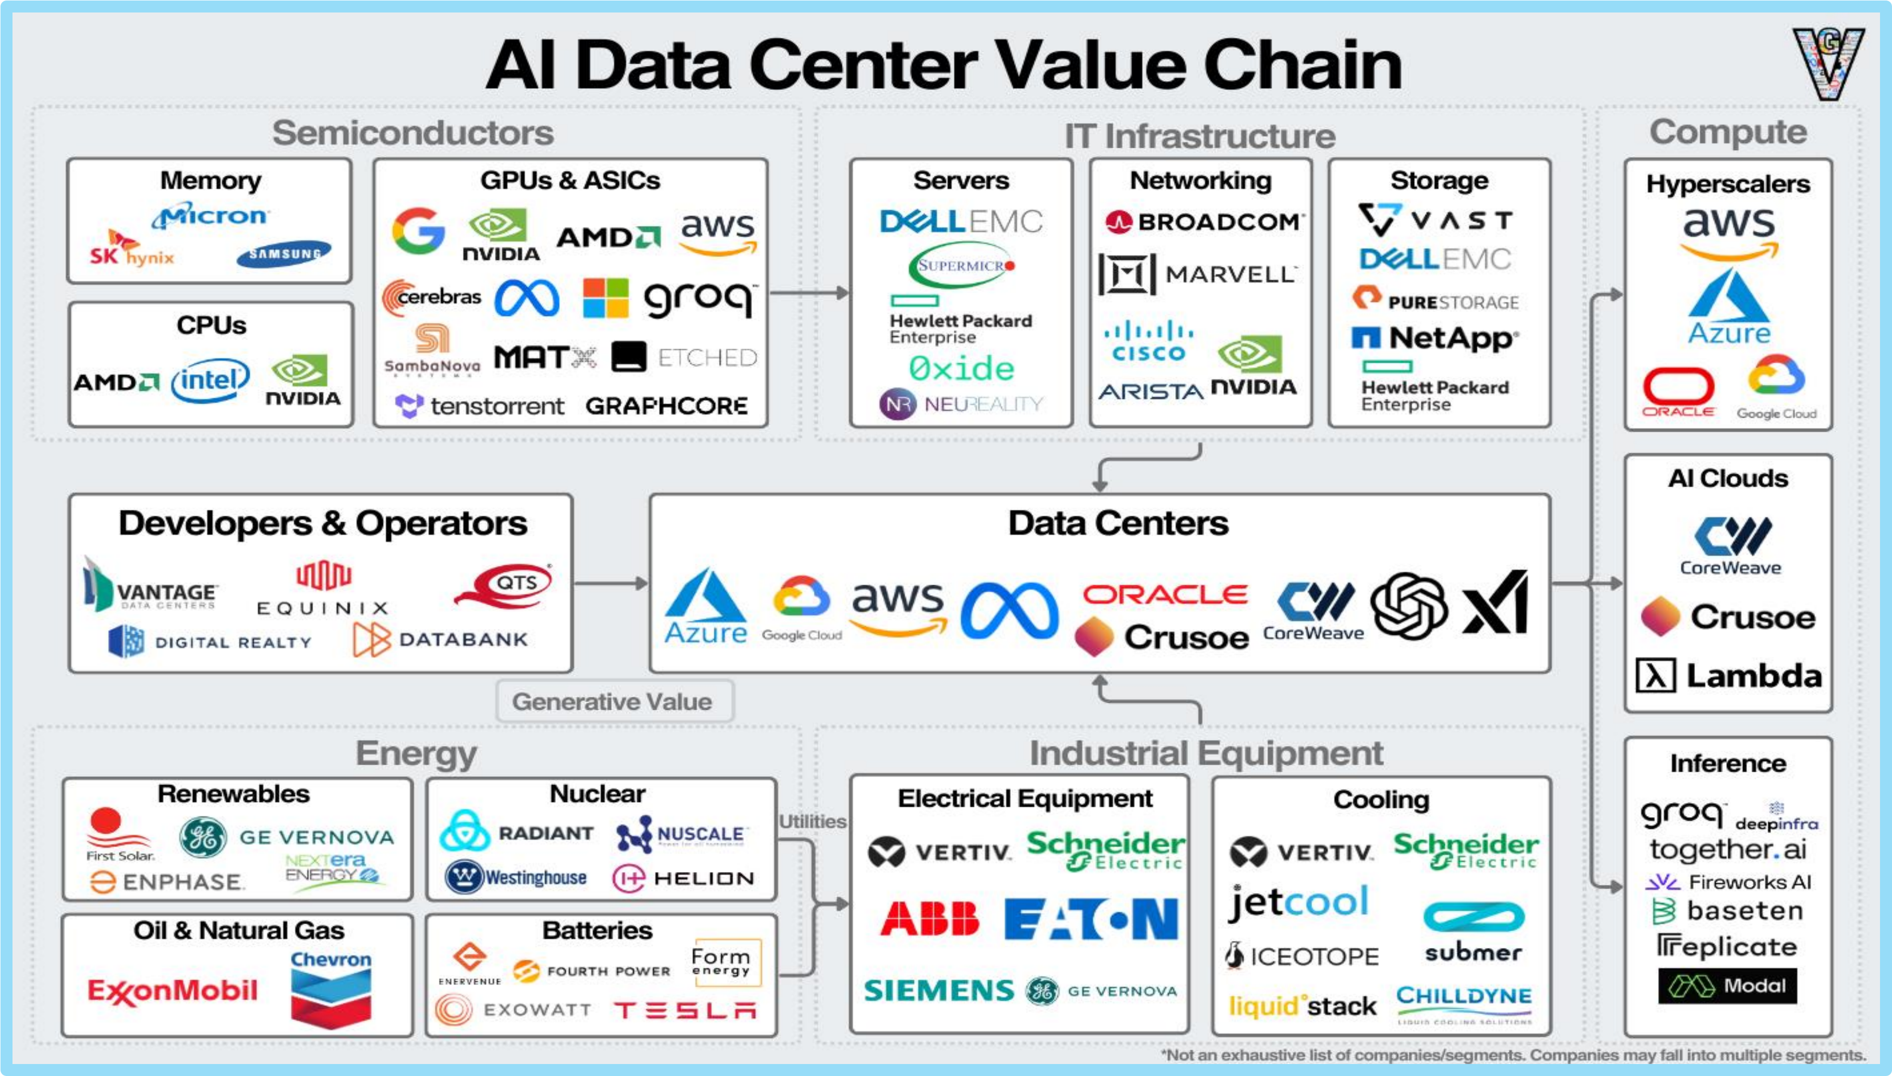

Source: Generative Value

CoreWeave

Increasing demand for AI solutions, is driving expansion for each core functionality of the value chain

Chip manufacturers:

Source of computational power required for training and inference of AI models

Key players:

NVIDIA, intel, AMD

Cloud service providers:

Backbone to the digital economy, by renting out storage, software, and computing power from data centers

Google Cloud, Microsoft Azure, Meta, CoreWeave, aws

AI companies

The builders of language models that are the foundation of AI solutions today

OpenAI, perplexity, deepseek, ANTHROP\C

CoreWeave

If the AI boom is the goldrush, then chip manufacturers are selling the shovels, and…

Chip manufacturers produce the **essential tool** needed to create AI models: chips that can run the calculations needed for training and inference

Market Leader: Nvidia

- Nvidia has 92% market share for GPUs
- Nvidia's success can be attributed to having a first mover advantage, originally pioneering GPU technology for gaming
- Since GPUs can process many calculations at the same time "in parallel" this is the perfect architecture for training and running LLMs that power generative AI, making GPUs an essential technology
- **Nvidia sets the industry standard for GPUs**

Semiconductor Industry Dynamics

- Fiercely competitive between industry rivals Nvidia, Intel, and AMD
- Chip manufacturers are facing the rise of application specific integrated circuits (ASICs): tailored chip architecture designed for efficiency and low cost
- Big Tech companies such as Google, Amazon, and Microsoft are designing these specialized chips in-house to avoid paying the Nvidia chip premium, and to reduce reliance

Sources: Yahoo Finance, Forbes, Financial Times, D&B Hoovers

The cloud providers are the gold mines and workers…

- The data centers run by cloud service providers represent the core of the AI environment with both chip manufacturers and AI companies dependent on data centers
- Chip manufacturers sell their infrastructure to data centers which act as a distributor, renting out computing power, storage, and software to customers
    - During the second quarter of fiscal year 2026, 88% of Nvidia's revenue was generated by data centers
- Data centers act as the home-base of the computing power needed for AI companies to train and run their models, making cloud service providers the integral piece to AI expansion and development

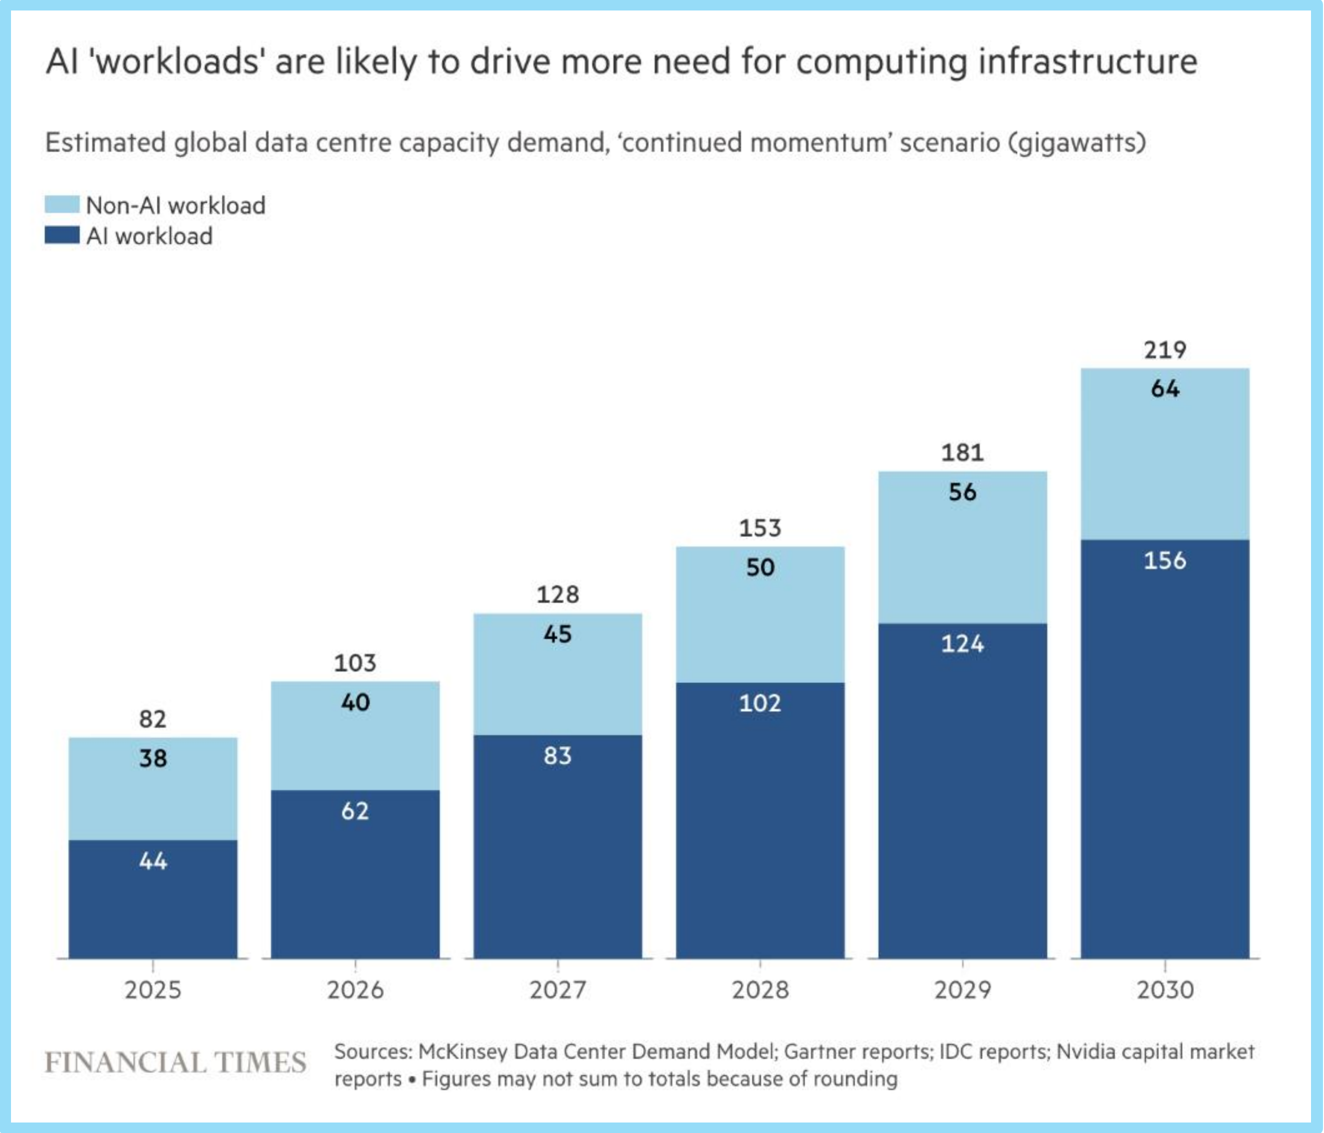

Sources: Yahoo Finance, Financial Times

CoreWeave

Leaving the AI companies' LLMs to be the gold

Large Language Models:
"Where traditional search engines and other programmed systems used algorithms to match keywords, LLMs capture deeper context, nuance and reasoning. LLMs, once trained, can adapt to many applications that involve interpreting text, like summarizing an article, debugging code or drafting a legal clause. When given agentic capabilities, LLMs can perform, with varying degrees of autonomy, various tasks that would otherwise be performed by humans (IBM)."

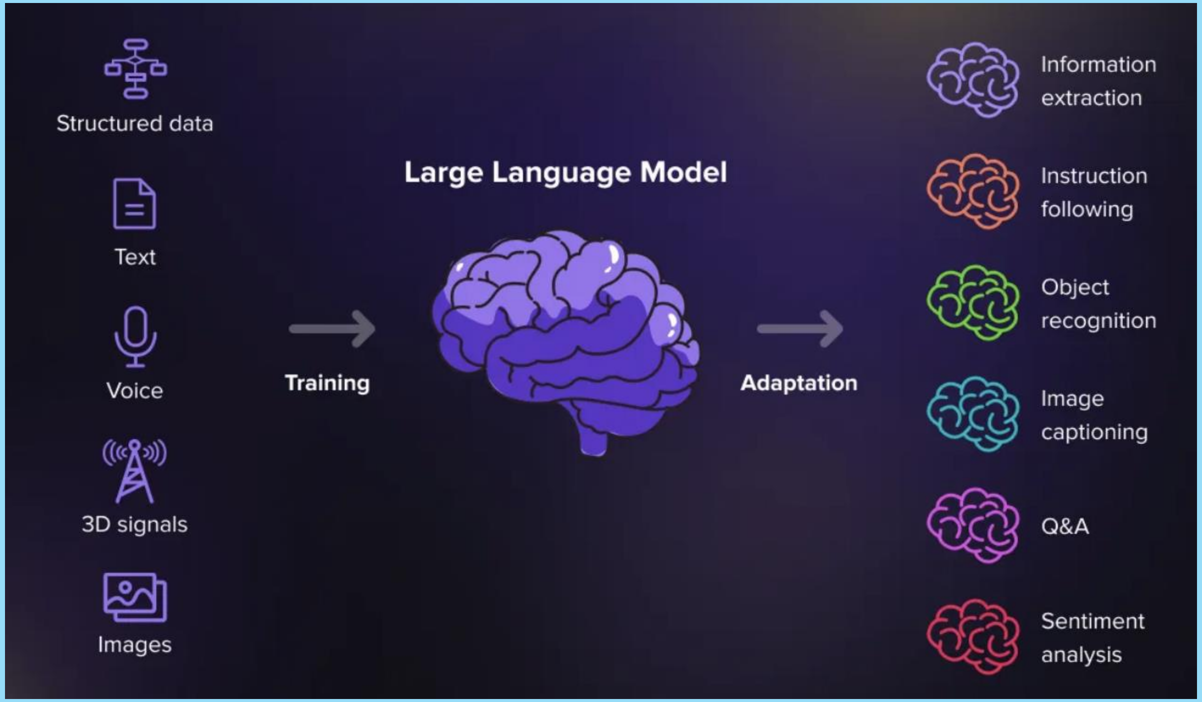

Large Language Models are accessible through interfaces such as…

Claude ChatGPT Gemini LLaMA by Meta

Source: IBM

CoreWeave

Exponential industry growth presents the opportunity to be the leader in delivering the AI infrastructure needed for the future demand of AI computing

Data centers are now a sought-after asset-class due to the increasing demand for AI solutions

- Traditionally, data centers were self-funded by cloud service providers, however given the rapid increase in demand for computing power, high levels of competition to be the leading supplier for AI, and scale of investment data centers require, half of today's infrastructure expansion is being financed by investors and developers through debt
- By 2030, the data center market is expected to be valued at $1 trillion, growing at a CAGR of 10%
- The demand for data center capacity could more than triple by 2030, with AI accounting for 70%

The key question is not who will capture this growth opportunity today, but rather who will secure it tomorrow and in the future

Source: Financial Times

CoreWeave

However, questions of a potential AI bubble could render AI infrastructure development obsolete

There are many risks associated with the AI industry, the main one being whether the level of investment in AI is causing a bubble

- The money spent on AI infrastructure this year has contributed more to GDP growth than consumer spending
- The rate and scale of investment in the data center expansion is causing concerns of overcapacity, future profitability, and the strain on energy demands
- AI companies are nowhere near bringing in the revenue necessary to justify the level of expenditure
- Given the rapid development of AI, it is uncertain whether current technology becomes obsolete, requiring new investments before infrastructure development is completed and ROI received

The challenge: Tech companies face the choice of either spending too much and wasting billions, or spending too little and being left out in the next revolutionary technological innovation

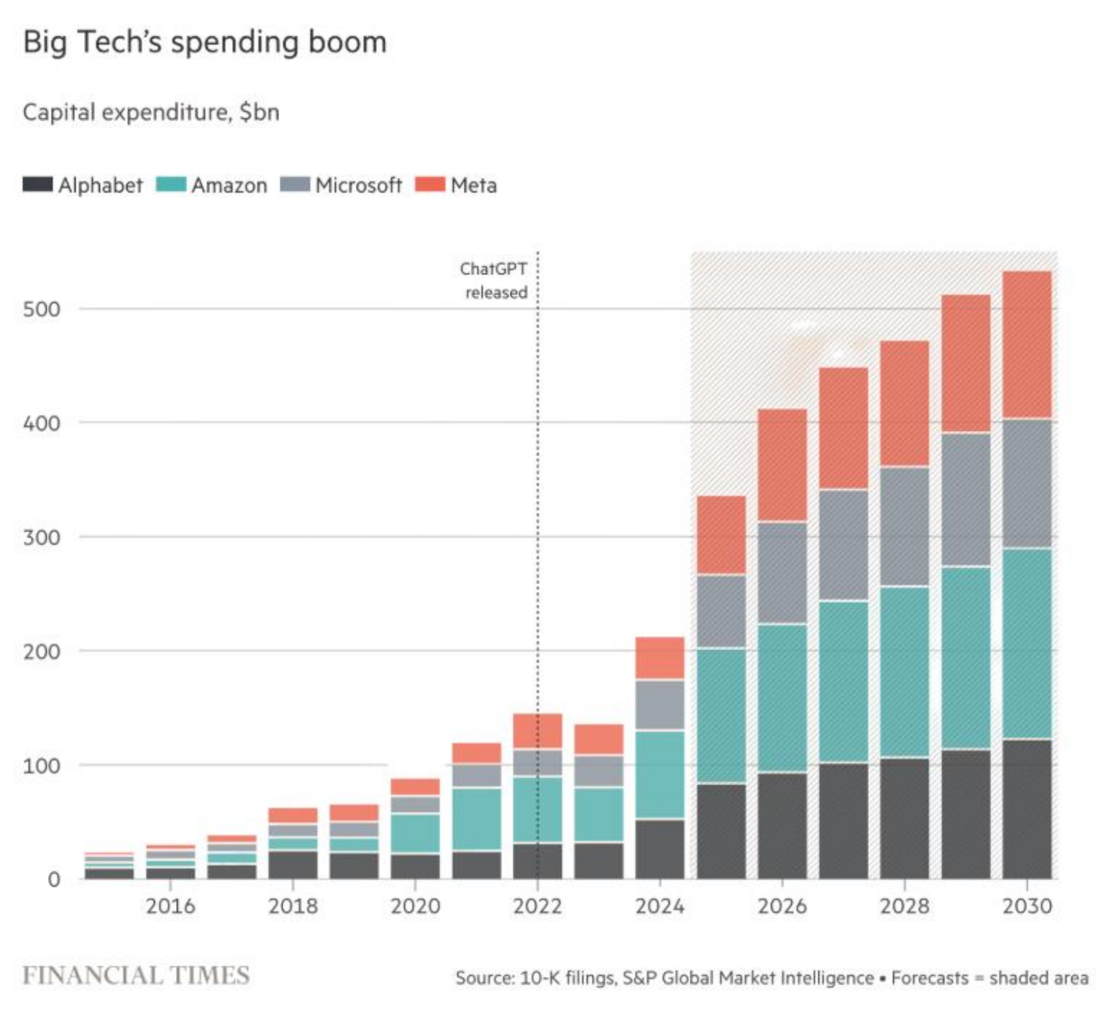

Source: Financial Times, Business Insider

CoreWeave

The next wave of AI is trending towards long-thinking models and agentic AI expansion, requiring more computing power than ever before

The rise of today's AI models are requiring increasingly more computing power:

- AI models are moving towards long-thinking, taking longer to process with the intention of reducing errors
- The industry is also seeing the rise of AI agents, models that can reason and act autonomously to adapt, plan, guide, and make decisions

The trends of long-thinking and AI agents signal that the future of AI will be more advanced and require more computational firepower to get there

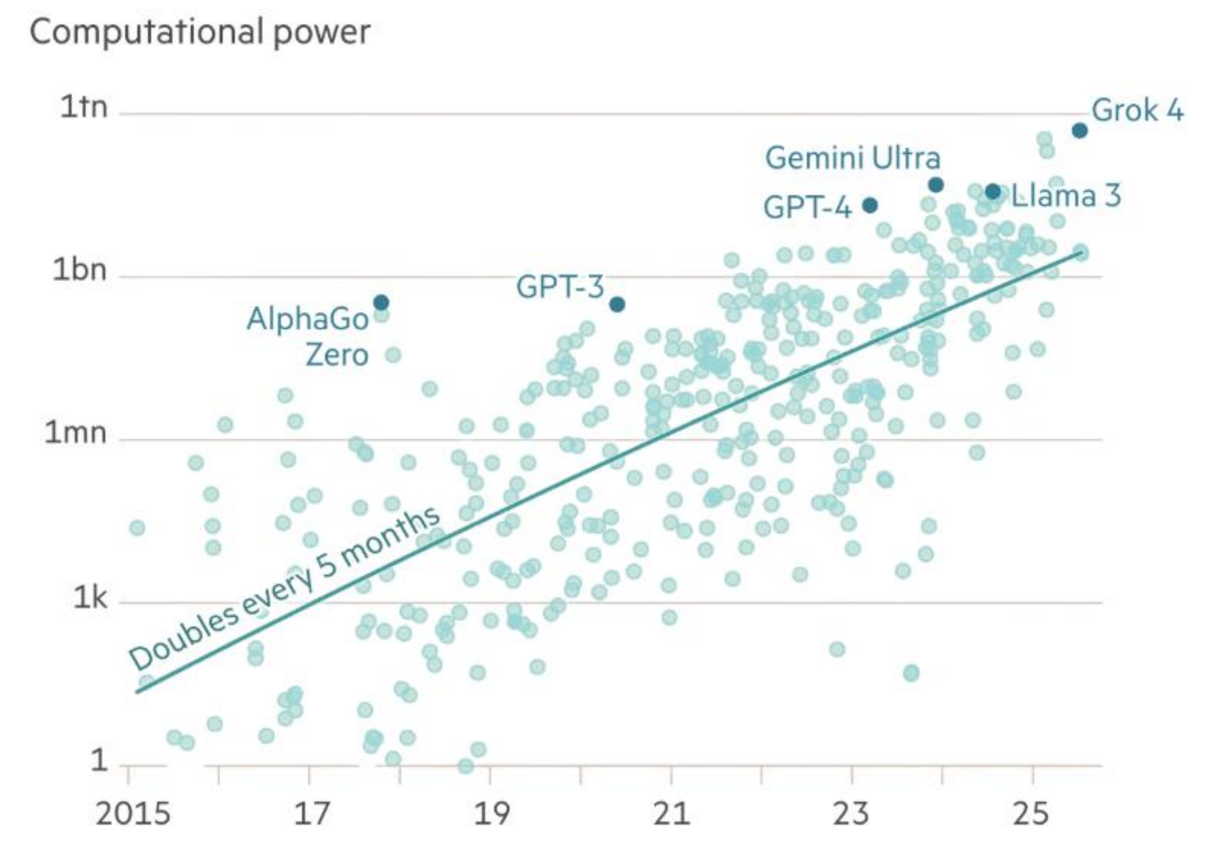

Sources: Wall Street Journal, McKinsey & Company, Financial Times

The future of AI will be decided by those who control and optimize the infrastructure to power the increasing demands for computing

The increasing demand for AI solutions require more from every step in the AI value chain: from computing, to storage capacity, to processing power, and to the energy required to run and cool equipment

The current approach to improving AI models is increasing the quantity of data needed to train language models, which will only continue the future demand for computing power provided by cloud service providers, through data centers

Data centers capture every element of this value chain making them the key determinant of who the winners and losers will be

Sources: Financial Times, Business Insider

Infrastructure Wins the Race

CoreWeave

The AI revolution isn’t happening in the cloud, it’s being built on the ground

The “cloud” is a marketing term: the real AI race is fought in megawatts, megabytes, and miles of copper wire

- Every AI function depends on physical infrastructure: servers, energy, cooling, and land
- Average data center ≈700 watts/GPUs ; one hyperscale site can require as much power as 80,000 homes
- As demand accelerates, controlling the infrastructure, not algorithms will determine who wins the AI race

Source: International Energy Agency

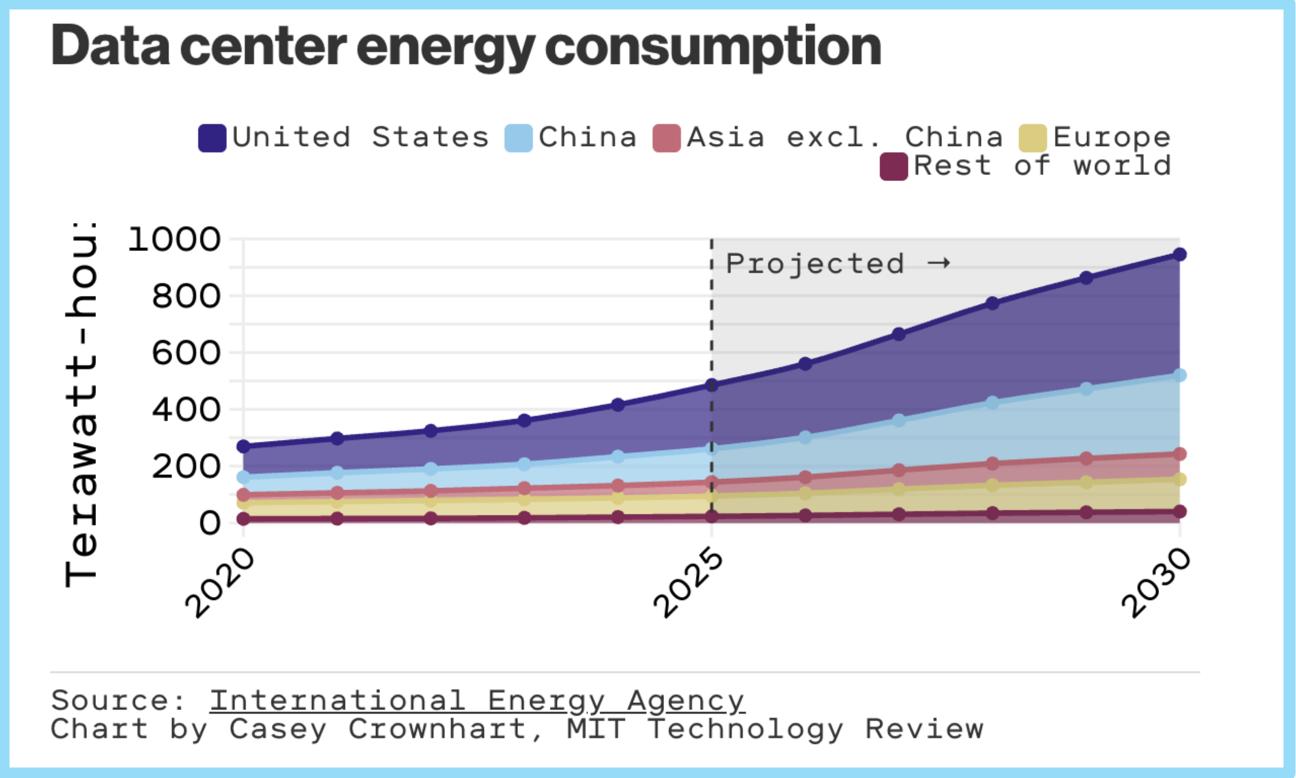

Global data center demand is set to surge from roughly 415 TWh today to 945 TWh by 2030 — equal to Japan’s total power use. Energy will define the limits of AI.

CoreWeave

02 Introduction to Data Centers

The global AI data center sector is both massive and rapidly evolving, serving every layer of the AI environment

$934B by 2030

- **Market Growth**: Valued at $236B in 2025, projected to reach $934B by 2030 *(31.6% CAGR)*
- **Core Role in AI Value Chain**: Provides the compute power behind LLMs, inference engines, and edge AI — the foundation for all model performance
- **Multi-Layer Integration**: Supports every layer, from chipmakers (NVIDIA, AMD) to AI companies (OpenAI, Anthropic) and cloud providers (AWS, Microsoft)
- **Investment Surge**: Increased investment from hyperscalers' capex cycles and sovereign AI infrastructure initiatives
- **Profit Center**: Data centers deliver higher long-term ROI than AI software itself, securing their position as the most prosperous link in the AI value chain

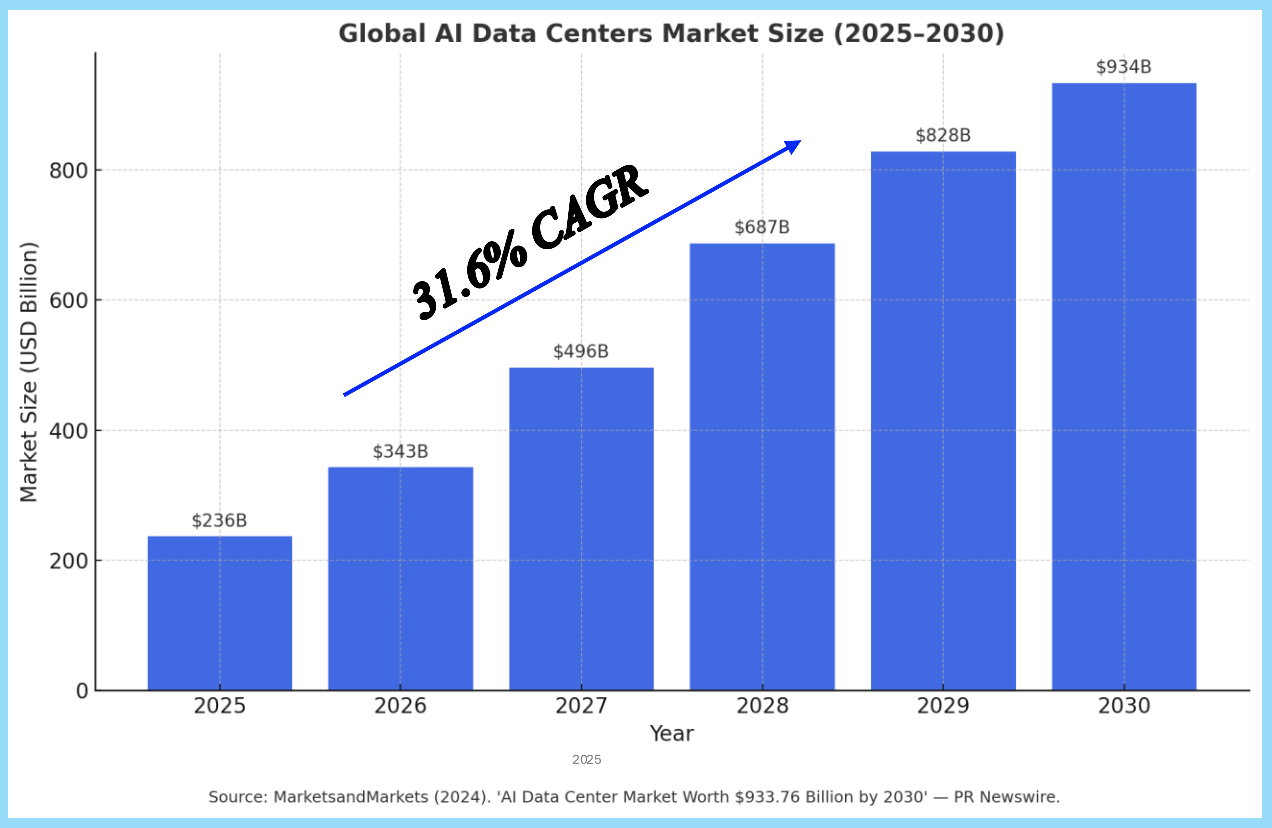

Source: MarketsandMarkets

Data centers are complex infrastructure hubs comprised of resource intensive computing units

- Compute, storage, and network interdependence define system limits
- Efficiency drops fast when any layer misaligns
- Scaling beyond 300 MW faces grid and heat constraints
- Millisecond latency directly impacts AI inference costs
- GPU and fiber shortages extend build cycles 18 + months
- Financing lags tech upgrades by up to a year

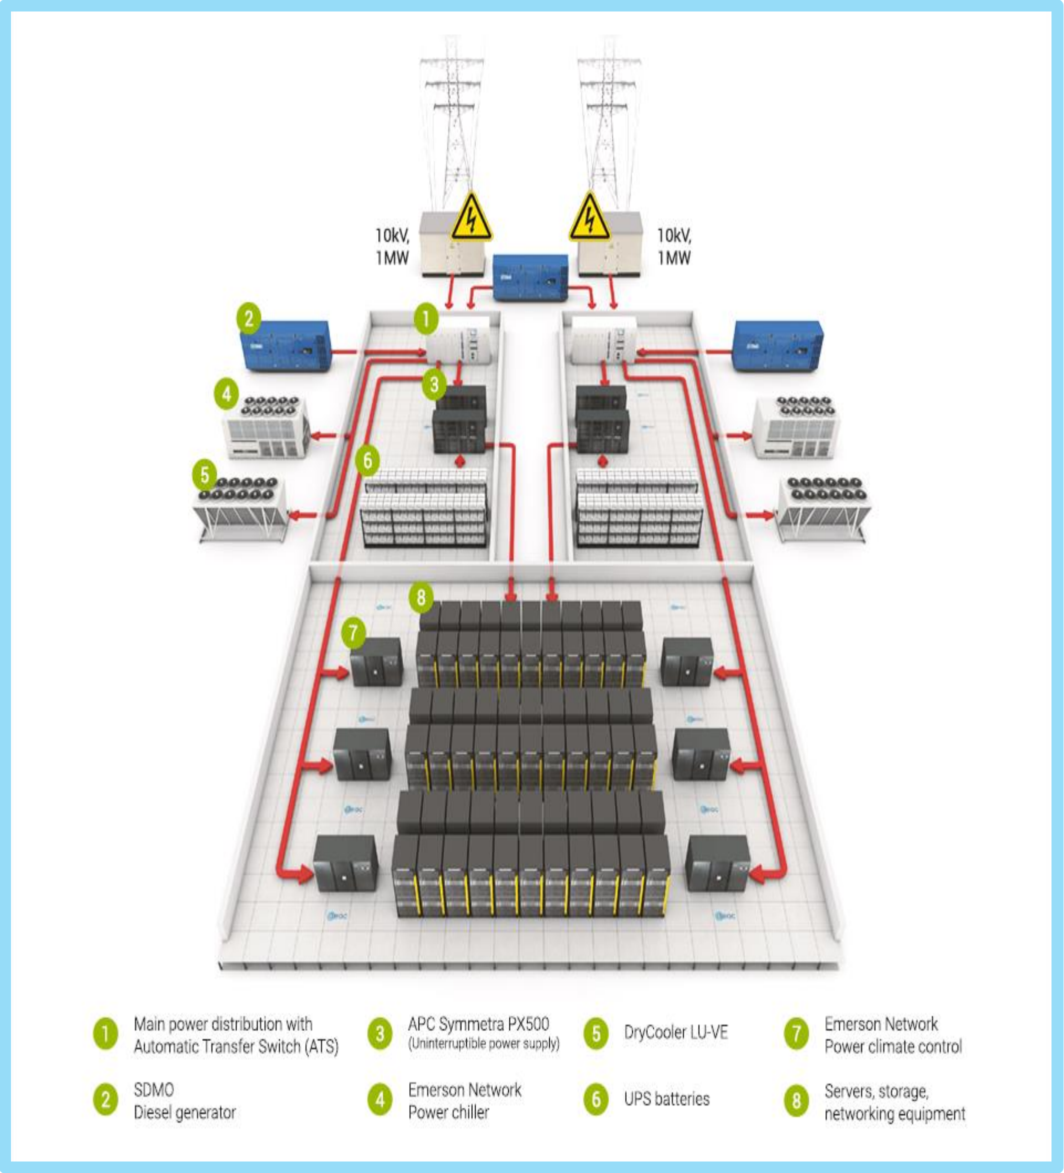

Source: International Energy Agency

Greenfield vs. Brownfield

Choosing between undeveloped expansion and adaptive reuse defines cost, speed, and sustainability in the AI data center race

Feature	Greenfield Project	Brownfield Project
Land Use and Development	Developing on previously undeveloped land	Redeveloping previously used sites
Costs and Timeframes	Higher upfront costs, faster implementation	Lower land acquisition costs, longer timelines, and potential expenses
Risk and Complexity	Lower risks, but must navigate regulations	Higher complexity due to contamination and environmental assessments
Environmental and Regulatory Considerations	Must comply with zoning and environmental impact assessments	Require rigorous remediation and compliance with specific regulations

Full control and long-term scale, but massive cost and slow ROI

Cheaper, faster, and politically cleaner path to compute power

Source: Harbinger Land

Computing units require a host of supporting assets to ensure proper usage and efficiency

Power and Electrical Systems

- Delivers electricity from the grid to all servers
- Includes transformers, UPS, and backup generators
- Keeps AI workloads running during surges or outages

Cooling and Environmental Control

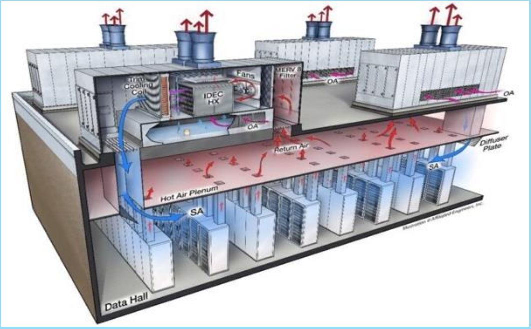

- Removes heat from dense, high-powered GPUs
- Uses chillers, airflow, and liquid cooling systems
- Prevents overheating and keeps performance stable

Racks, Servers & Hardware

- Metal frames housing CPUs, GPUs, and storage
- Provide structure, airflow, and organized cabling
- Enable easy scaling and efficient compute power

Network and Connectivity

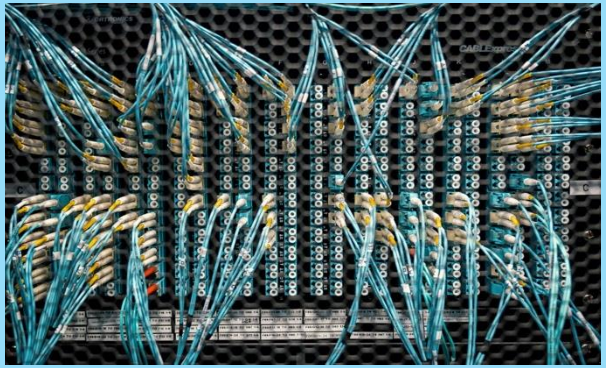

- High-speed fiber links servers and the internet
- Routers and switches move data with low latency
- Ensures fast AI training and reliable communication

Monitoring & Security

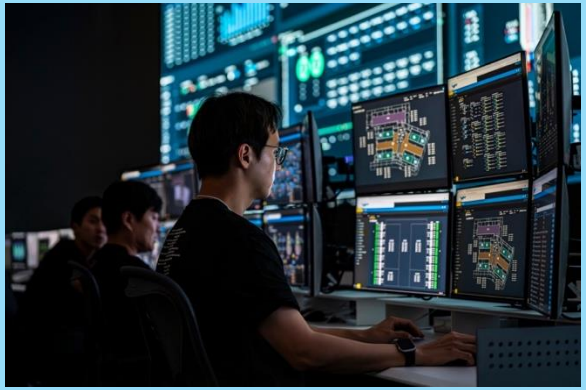

- Cameras, sensors, and biometric access control
- Tracks power, temperature, and data activity
- Protects uptime, safety, and sensitive information

Source: Data Center Dynamics

With the resource intensive nature of data centers, concerns of environmental impact have arisen

- AI data centers on pace to emit about 340 Mt of CO_2 annually by 2030
- About 50% of U.S. data centers are in areas with moderate to extreme water stress (Arizona, Texas, and Nevada)
  - Each hyperscale facility evaporates millions of gallons of freshwater daily
- Cooling discharges and thermal pollution threaten local ecosystems and groundwater quality
- Renewable offsets don't fully neutralize the carbon and material footprint of steel, concrete, and semiconductor supply chains
- Concentration increases competition for limited water resources as AI-driven cooling demand grows
- **Sustainability now defines scalability, as unchecked expansion risks trading innovation for ecological instability**

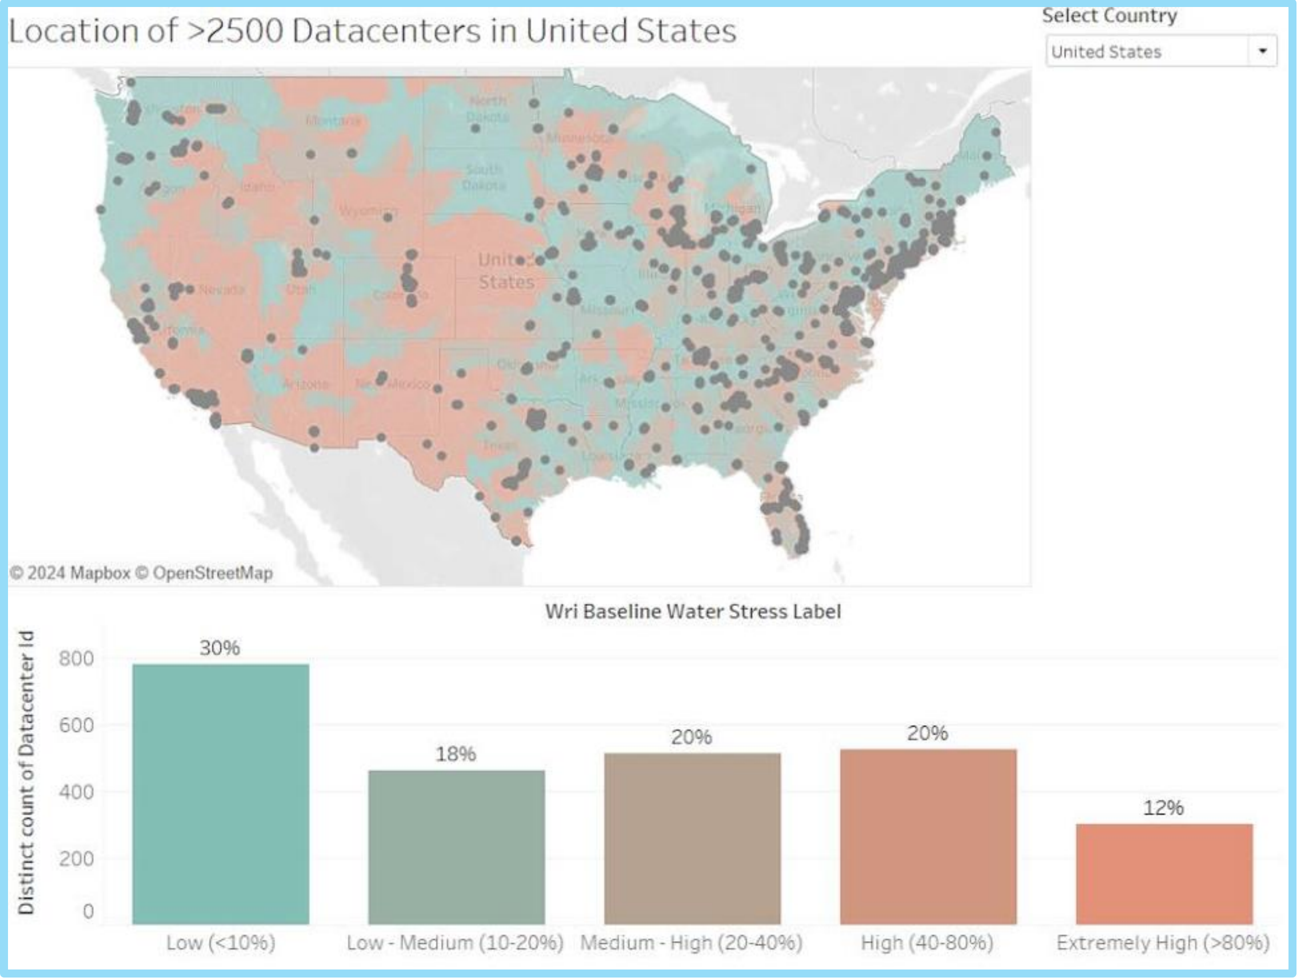

Source: Planet Tracker & IEA

CoreWeave

Sustainability isn't optional, it's the next infrastructure constraint

- Data centers now produce e-waste and embodied carbon equal to 3–4% of global IT emissions
- Cooling and power-delivery systems make up for 40% of total energy draw
  - Newer liquid-cooling and immersion systems are cutting this by half in pilot sites
- Materials and construction (steel, concrete, servers) represent the hidden emissions frontier; many "green" centers ignore this life-cycle footprint, these are known as Scope 3 emissions
- They remain the largest unsolved sustainability challenge for AI hyperscalers:
  - Make up 70–90% of total carbon output
  - Green electricity can reduce operational impact but can't offset the embodied footprint of the infrastructure itself
- The next "winner" won't just build faster: they'll build smarter, circular, and self-sustaining infrastructure

➡ **Possible early models of circular infrastructure: waste-heat reuse projects in Denmark, Sweden, and Oregon heating neighborhoods**

Source: Net Zero Insights

CoreWeave

Today's hyperscalers' rapid investment in data centers pose ethical concerns to neighboring communities

- Data centers spark short-term construction booms but create only about 0.2 permanent jobs per megawatt once operational
- Utility load growth from new data centers (30–45% by 2032 in Kentucky) leads to higher electricity costs for residents and small businesses
- Housing pressures rise as industrial zoning replaces affordable residential areas, as a result housing rates increase and displace long-term residents
- Residents face constant noise, vibration, and heat output from 24/7 cooling systems, lowering the livability and local air quality
- Water use conflicts are rising: one hyperscale facility can consume millions of gallons daily, competing with nearby towns
- Companies continue to advertise "green AI," while locals often see industrialization without inclusion
  - Limited job creation and rising household costs

➡ **Hyperscalers strain communities, driving up housing costs, water and energy usage, and pollution while leaving limited lasting value for locals**

Source: PWC and LG&E Energy

CoreWeave

Unchecked expansion turns local growth into social backlash

1. Unemployment Issues:

- After construction, data centers employ less than one full-time worker per 5 MW of capacity, creating minimal long-term opportunity
- Most roles demand advanced engineering or IT skills, leaving legacy workers from manufacturing or trades excluded from new growth
- This gap fuels local resentment—residents power the infrastructure but aren't hired to sustain it

2. Economic Inequality:

- Housing and living costs spike as high-income technical employees move in, pricing out long-term residents
- Local wages lag behind corporate profits; economic benefits concentrate in metropolitan HQs rather than host towns
- Tax incentives meant to attract hyperscalers often erode local budgets, reducing funds for schools, roads, and community services

3. Public Health and Quality of Life:

- Noise, diesel exhaust, and heat pollution from 24/7 operations increase respiratory problems and stress levels
- Light pollution and vibration from cooling systems disrupt sleep, wildlife patterns, and mental health
- Residents living near hyperscale clusters report declines in property satisfaction and rising energy costs

4. Community and Civic Impact:

- Zoning disputes, broken job promises, and limited transparency fuel distrust toward both companies and local government
- Cultural displacement transforms rural or suburban towns into industrial tech corridors, fracturing community identity
- Social fatigue sets in as citizens feel excluded from planning processes while bearing higher costs and fewer benefits

Source: Tech Policy Press & Business Insider

CoreWeave

03 Analysis of Winners and Losers

CoreWeave

With ethical concerns and project bottlenecks, some players have won big while others have lost large sums attempting to keep up with computing infrastructure demand

Winners: Efficient Builders and Strategic Partners	Losers: Overextended or Misaligned Developers
Aligned early with GPU suppliers (NVIDIA, AMD) to secure capacity and favorable pricing	Overcommitted to greenfield construction, facing long delays and environmental resistance
Adopted brownfield expansion models → faster and lower-risk builds	Underestimated cooling and energy costs →budget overruns and idle assets
Integrated sustainability innovations: liquid cooling and renewable power sourcing to pass regulatory hurdles	Failed to anticipate GPU shortages → underutilized or incomplete facilities
Partnered with hyperscalers with leasing or joint-venture models → spread capital risk	Ignored local pushback → project cancellations or expensive redesigns
Result: Shorter deployment cycles, lower capex per MW, and faster go-to-market capacity	**Result: Billions in stranded capital and missed AI infrastructure opportunities**

Source: McKinsey & Company

Winners aren't just keeping up with AI demand, they're redefining what efficient, sustainable, and scalable data center construction looks like

aws

Northern Virginia

- Largest AI cloud hub in the U.S. with hundreds of data halls built in parallel
- Prefabricated infrastructure approach reduced build time per hall by 35%
- Integrated custom Trainium and Inferentia chips to optimize energy per training cycle

Mastered scale and efficiency through vertical integration

CoreWeave

New Jersey

- Repurposed industrial warehouses to deploy GPU clusters 3x faster than new builds
- Focused on AI leasing model—providing instant access to compute capacity for clients like OpenAI
- Backed by NVIDIA partnerships and $7B+ financing, enabling rapid national expansion

Turned brownfield agility into a market-leading advantage

Google Cloud

Hamin, Finland

- Converted an old paper mill into a state-of-the-art AI data center (brownfield build)
- Utilizes seawater cooling from the Gulf of Finland—cutting cooling energy use by ~40%
- Fully powered by renewable electricity, enabling true carbon neutrality

Proves efficiency and sustainability can coexist at global AI scale

Microsoft

Des Moines, Iowa

- $3.5B AI campus expansion designed specifically for high-density GPU workloads
- Custom liquid cooling systems boost compute efficiency by 20–30%
- Built using modular data hall design, allowing phased deployment at record speed
- Renewable energy partnerships secured grid capacity and fast-tracked permits

A benchmark for rapid, sustainable hyperscale construction

Source: Data Center Frontier & Data Center Dynamics

CoreWeave

Colossus is the peak of brownfield expansion, but could also be the breaking point

- xAI’s Colossus: converted former Electrolux factory in Memphis into a hyperscale AI facility in ≈122 days, fastest brownfield redevelopments on record
- Built to support 100,000+ NVIDIA H100 GPUs
  - Reusing existing grid links, cooling systems, and heavy structural loads drastically shorten deployment timelines
- Shows efficiency advantage of brownfield sites
  - Massive resource concentration required to achieve it
- Benefits from private energy resources, concentrated capital, and direct executive control
- Rapid scale-up raises concerns over grid strain and environmental impact
  - Emphasizes operational trade-offs behind extreme growth

Brownfield conversions are the most viable path for expanding AI infrastructure, but speed without sustainability introduces long-term risks

Source: xAI

CoreWeave

Losers have failed to construct greenfield sites as a result of environmental concerns, zoning faults, and resource shortages

aws

Ireland

- Expansion plans halted by national energy grid moratoriums on new data centers
- Government cited severe power shortages and strain on Dublin's infrastructure
- AWS delayed multiple facilities as energy allocations were capped nationwide

Regulatory limits and resource scarcity have slowed even the largest cloud provider

Meta

Odense, Denmark

- 2.8 million sq ft expansion canceled mid-construction due to cost inflation
- Public backlash grew over energy consumption and minimal local benefit
- Meta shifted toward AI-specific retrofits instead of new greenfield sites

Sustainability pressure and cost volatility derailed Meta's growth strategy

Google Cloud

The Dalles, Oregon

- Faced multi-year legal disputes over groundwater rights and secrecy
- Drought conditions triggered public protests and environmental review
- Local courts restricted future water allocations, capping expansion potential

Community and environmental resistance now define Google's build limits

Microsoft

Holland Township, Michigan

- Proposed $1 billion AI campus met intense zoning and farmland opposition
- Residents cited noise, visual, and infrastructure impacts as key concerns
- County revoked permits, sending the project back to review

Local politics and social pushback can outweigh corporate scale

Even hyperscalers are struggling to overcome grid limits, regulations, and community opposition. For CoreWeave to continue to implement its strategy of high speed of deployment and low upfront investment, it must leverage brownfields to stay adaptable in the next phase of AI infrastructure.

Source: Data Center Frontier & Data Center Dynamics

CoreWeave

Building AI infrastructure is a high-stakes bet on capital, time, and uncertainty

- Global AI data-center investment projected to exceed $6.7 trillion by 2030
- Projects require $500 million–$2 billion per site with 18–36 month build cycles
- Rising costs for GPUs, cooling, and power grids strain returns as financing lags demand
- Market saturation and permitting delays increase stranded-asset risk
- Investors face growing pressure as ROI timelines stretch beyond typical tech cycles
- Capital-heavy builds risk locking firms into outdated designs before payback is achieved

The next infrastructure winners will scale capital intelligently—not endlessly

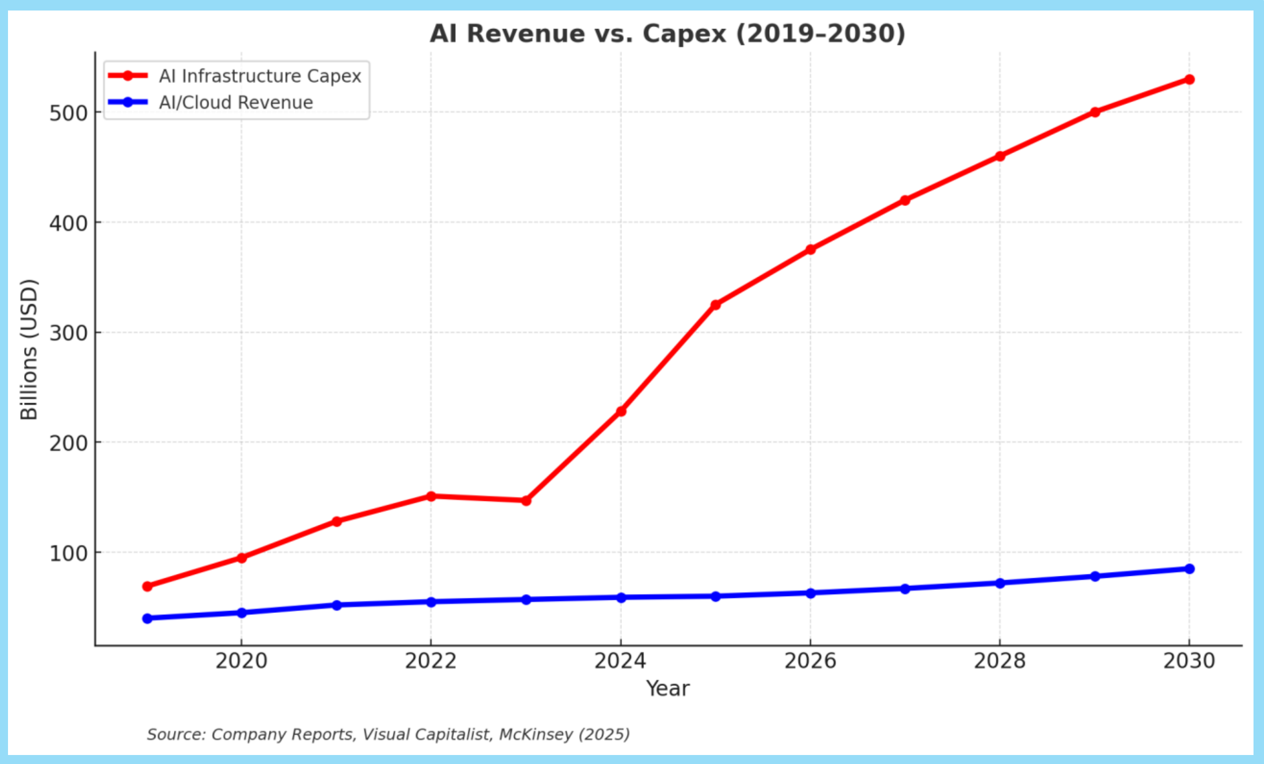

AI infrastructure spending is rising faster than cloud revenue growth — signaling a looming imbalance between cost and return.

Source: McKinsey

CoreWeave

Mitigating risk can be a challenge, but CoreWeave has brought a new lower risk strategy to the space

CoreWeave's Strategic Advantage:

- CoreWeave (CW) operates a specialized GPU cloud, built for AI, visual effects, and high performance computing workloads
- Instead of owning hyperscale campuses, CW leases capacity from existing providers and installs GPUs to resell compute power
- CoreWeave's vertically integrated stack spans GPU orchestration, data center leasing, and managed AI services

Differentiating Impacts

- This lowers CapEx exposure and speeds time-to-market, avoiding the multibillion-dollar build cycles that burden traditional cloud firms
- Backed by NVIDIA and Magnetar Capital, CW's flexible model positions it as the go-to GPU supplier for AI startups and enterprises
- As hyperscalers chase scale, CW advantage lies in specialization, adaptability, and speed

Source: Coreweave and Data Center Dynamics

CoreWeave

CoreWeave is a leading leasing agent of AI computing power with explosive growth and extensive demand for services

- **Explosive Growth:** CW revenue surged 207% YoY in 2025, fueled by hyperscaler GPU shortages and AI model demand
- **Locked in Demand**: Over $30.1 B in signed orders secures years of guaranteed compute leasing contracts
- **Risk Mitigation**: The $6.3 B NVIDIA capacity guarantee (through 2032) secures buybacks of unused GPU supply, insulating CW from demand volatility and enabling sustained scaling without overbuild risk
- **Asset-Light Strategy**: Instead of owning mega campuses, CW leases and retrofits sites, converting fixed CapEx into scalable OpEx
- **Strategic Position**: Now acts as the "AWS for GPUs," bridging hyperscaler capacity with AI-native companies

CoreWeave's model transforms AI infrastructure from a build race into a leasing ecosystem

Source: Yahoo Finance, Financial Times, and Coreweave

Leasing is the main facet of business for CoreWeave, but the company leaves room for a powerful industry dominating expansion

- **Vertical Capture:** Retrofitted brownfield sites atop leased infrastructure secure power, control, and faster deployment
- **Distributed Scale:** Regional 100–300 MW clusters reduce latency, grid risk, and meet localized AI demand
- **Capacity Marketplace:** Enables firms to buy and sell GPU power under flexible SLAs and performance contracts
- **Enterprise & Government Reach:** Targets regulated and sovereign AI clients requiring data residency and transparency
- **Financing Engine:** Long-term take-or-pay contracts and renewable PPAs stabilize cash flows and lower WACC

CoreWeave's expansion integrates GPUs, hybrid infrastructure, AI marketplaces, and strategic partnerships into a vertically aligned model that scales faster, cheaper, and smarter than traditional hyperscalers.

Source: Bain & Company, Barron's and Coreweave

From GPU orchestration to financing, CoreWeave connects every layer of the AI stack

GPU Orchestration Layer

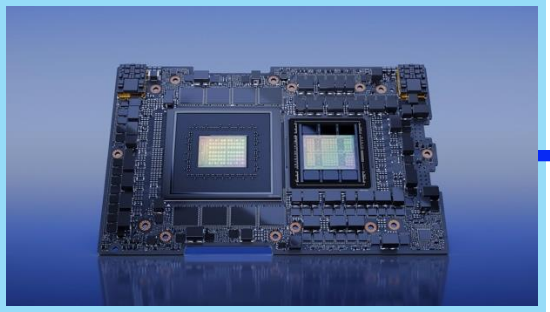

CoreWeave installs and manages NVIDIA GPUs across leased and owned sites.

Hybrid Infrastructure Expansion

Retrofitted brownfield facilities enable faster, lower-cost scaling across regional clusters.

AI Compute Marketplace

Platform for trading and allocating compute power under SLAs.

Strategic Partnerships & Financing

NVIDIA backing, long-term PPAs, and take-or-pay contracts stabilize growth.

Source: Bain & Company, Barron's and Coreweave

CoreWeave

04 Recommendation

CoreWeave

CoreWeave's current strategy is leasing infrastructure from large providers in order to install GPUs to sell computing power to other firms

GPU Leasing Engine
CoreWeave installs NVIDIA GPUs across leased hyperscaler facilities, then sells compute power to AI clients under flexible service contracts

Speed to Market
Using pre-built data center capacity allows deployment in weeks instead of years, meeting surging AI demand immediately

Capital Efficiency
Avoids the $500M–$2B per-site cost of new data centers, converting fixed CapEx into variable OpEx for faster payback

Transition Risk:
Continued hyperscaler dependency limits control over energy sourcing, site optimization, and long-term margin capture

Strategic Dependence
Reliance on hyperscaler space (AWS, Google, Microsoft) exposes CoreWeave to pricing power and supply constraints

Source: Bain & Company, Barron's and Coreweave

CoreWeave

CoreWeave's future scalability is overly dependent on leasing current hyperscaler infrastructure, putting CoreWeave in a vulnerable position

CoreWeave relies heavily on rented capacity from larger operators. This approach supports fast deployment today, but it limits control of long-term scale. As demand for computing grows, leased capacity becomes less predictable and more expensive. Without more owned sites, CoreWeave risks capacity shortages when it needs them most.

- Lease availability is shaped by the priorities of the provider, not CoreWeave
- Rising demand makes long-term space harder to secure
- Competing tenants restrict power and space allocations
- Limited influence over facility layouts, cooling plans, and future upgrades
- At will of leading party ownership even with extensive investing in CoreWeave GPU technology within leased centers

Key Challenge: Rising lease prices and tightening vacancy increase CoreWeave's risk if dependence on external providers continues.

Source: Bain & Company, Barron's and Coreweave

CoreWeave

From dependency to control: CoreWeave's cost efficient, rapid, and scalable answer to the AI wars

Question

How can CoreWeave scale their capacity and limit dependence on hyperscalers while simultaneously maintaining a fast deployment strategy?

Answer

Revitalize and retrofit a brownfield industrial site into a state of the art AI data center for lower cost, allowing CoreWeave to install capacity in a fraction of the time required for new construction. This aligns with CoreWeave's preference for rapid deployment, while also giving the company greater control over its long-term infrastructure plan.

CoreWeave

CoreWeave must rapidly and efficiently build its own data centers through brownfield sites to achieve hyperscaler independence

Retrofit brownfield industrial sites into data centers to make CoreWeave more independent from hyperscalers

Develop strategic autonomy from hyperscalers

Optimize cost and speed

Scale to vertically integrate

Ideal site selection and quick deployment

Leverage partnerships with infrastructure providers

Expand compute capacity

CoreWeave

CoreWeave needs to strengthen its hybrid model through their own data center expansion

- Leasing alone cannot guarantee future access to the volume of racks required
- Owning sites creates a dependable base of power and floor space
- A hybrid model protects CoreWeave from supply shortages
- Site ownership supports multiyear contracts with customers
- Leasing does not guarantee efficient computing conditions (stable electric, cold water)
- As prices rise to lease data center space, leasing is becoming a risky position (~20% increase YOY)

Key Takeaway: Continuing to rely on external providers preserves speed, but it limits long-term planning. Expanding its owned footprint supports stability and cost control. A combined model provides the speed CoreWeave values while building a foundation for long-range capacity needs.

Source: CBRE

CoreWeave

Expanding computing capacity gives needed support to more intensive models while also keeping up with industry growth

The Need for Expansion:

- Wait times for new capacity in primary markets often reach 12 to 24 months due to low vacancy
- Asking prices in major markets rose 7 to 15 percent year over year

The Challenge

- Utility and infrastructure construction can represent 25 to 40 percent of total build cost in new sites

Our Strategic Answer

- Brownfield development avoids most of this cost and shortens build time
- Faster deployment lets CoreWeave secure customer demand before competitors add capacity
- Lower capital cost supports more flexible pricing and deeper market penetration
- More infrastructure allows for greater compute capacity

Key Takeaway: Speed of deployment is a competitive advantage because customers often require immediate access to computing capacity. Additional infrastructure expands computing capacity. Maintaining this pace at a reasonable cost strengthens CoreWeave's position. Brownfield locations make this possible by reducing the cost and time required to open new sites.

Source: Yahoo Finance

CoreWeave

CoreWeave must break away from hyperscaler dependence by owning the infrastructure stack that powers the next wave of AI

- **CW currently runs 250,000+ GPUs, over 33 data centers and roughly 1.6 GW of contracted power capacity**
- Through its acquisition of Core Scientific in July, CW gained ~1.3 GW of gross power capacity and remove ~$10 billion in cumulative lease obligations
- Hyperscaler reliance limits control over cost, capacity, and deployment speed.
- Leasing space from hyperscalers exposes CW to rising rental rates and delayed access to critical GPU and power infrastructure
- Owning data centers enables direct control over power, cooling, and networking the key inputs that determine compute efficiency and uptime
- Vertical integration reduces exposure to third-party risk and ensures predictable, scalable growth as AI workloads expand
- By developing its own sites, CW captures infrastructure value instead of paying long-term lease premiums to hyperscalers.

Key Takeaway: By owning both its assets and operations, CoreWeave positions itself as an independent infrastructure leader. It is purpose-built to support 100,000+ GPU megaclusters and 100s of MW of power, overcoming the bottlenecks that legacy data-centres face.

Sources: Goldman Sachs, CBRE, Data Center Dynamics

CoreWeave

Constructing a data center based on a brownfield low-cost model can assist CoreWeave in leveraging a lower risk position while also establishing CoreWeave as a key AI computing player

Example Site Overview:

- 1,200-acre former Kodak industrial campus in Rochester, NY
- 16M sq ft industrial space; 100+ tech tenants
- Heavy-industrial zoning continuously in place since the 1890s
- 100-acre contiguous parcel available-flat, rail-adjacent, next to 150 MW power and 1.5 MGD cooling

Example Site Key Existing Infrastructure:

- 50-100 MW CHP, expandable to 500 MW
- Chilled-water + steam networks built for film processing
- CSX rail spur with 20-bay docks
- Lake Ontario cold-water intake (4-6°C) with active permit
- Brownfield approach cuts 40-50% of typical capital vs. greenfield builds
- No zoning delays, no new environmental review, minimal new utility construction
- Retrofit 785,000 sq ft for rapid time-to-revenue

Strategic fit for CoreWeave: By utilizing this brownfield site with established utilities, zoning, and logistics, CoreWeave can bring new capacity online within months instead of years, turning underused industrial space into high-value AI assets

Source: Kodak Eastman

CoreWeave

CoreWeave aims to grow its relationship with Nvidia while leveraging a vertically integrated future

Strength of Partnership:

- One of CoreWeave's greatest competitive advantages is the strength of the company's partnership with Nvidia which produces the essential market standard GPUs
- CoreWeave is a preferred partner, acquiring thousands of Nvidia GPUs, and claiming to be the largest private operator of Nvidia GPUs in the US, with over 45,000 chips
- CoreWeave was the first to offer the complete BlackWell GPU portfolio at scale, positioning CoreWeave as the first readily available source for AI workloads
- Nvidia owns a meaningful equity stake in CoreWeave (various sources cite about ~6.6 % at one point) and is both supplier (chips) and key customer/partner

Recent Growth:

- CoreWeave is already further leveraging its relationship with Nvidia by expanding its operations in the UK through a $1.3 billion dollar investment deal to deploy chips in two new data centers
- In September 2025, Nvidia and CoreWeave signed a $6.3 billion cloud capacity agreement
- Nvidia will guarantee to purchase any of CoreWeave's unsold cloud-computing capacity under this agreement through April 13, 2032

Sources: CoreWeave, Financial Times, NASDAQ, Reuters

CoreWeave will gain greater compute capacity through infrastructure build-out

- In the U.S. there are 450,000 industrial brownfield sites
- Since 1995, the EPA has distributed $2.37B in grants to clean up brownfield sites for reuse
- CoreWeave has ~590MW of current power capacity with plans to reach 850MW by end of year
- CoreWeave has a backlog of $55.6B in potential revenue incentivizing CoreWeave to expand compute capacity

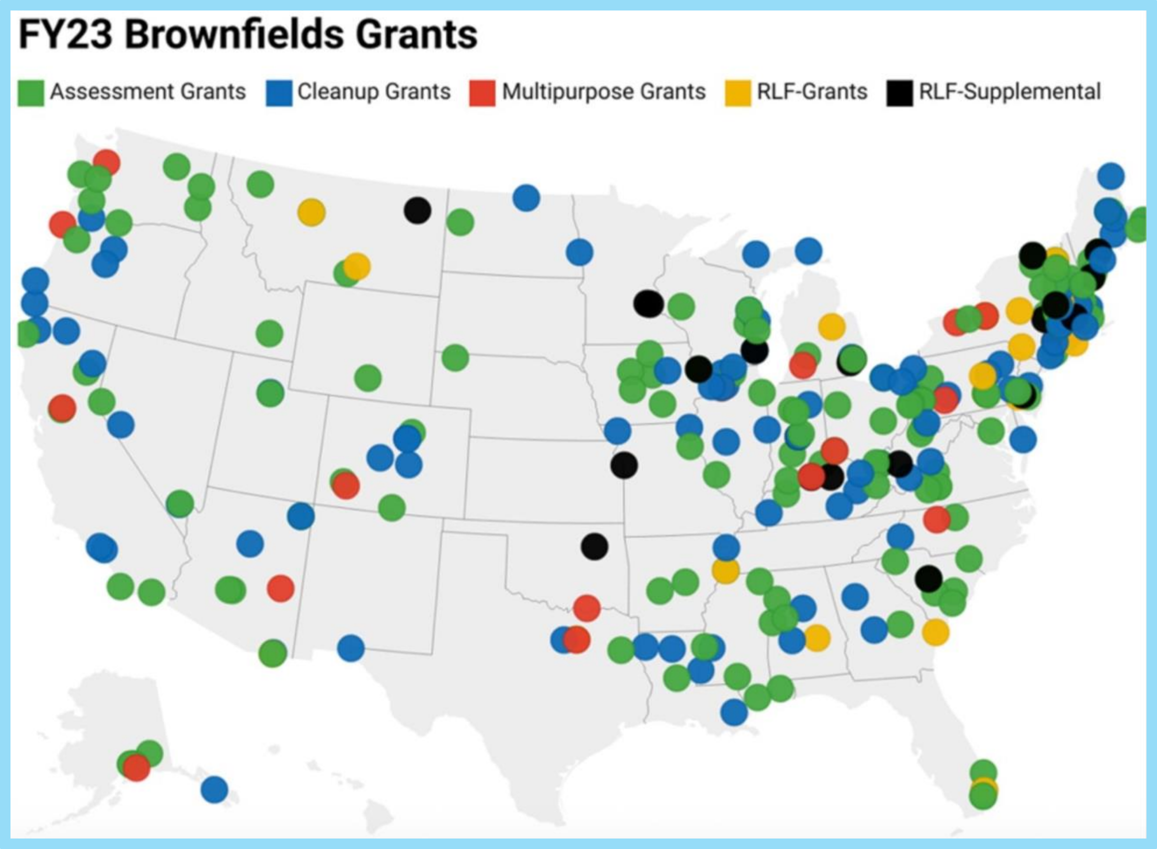

Source: EPA, AI Invest

In conclusion, CoreWeave must leverage a lower risk position by vertically integrating

- **Market:** Growth depends on infrastructure, not algorithms
- **Value Chain:** Chips + Cloud + Data Centers = AI backbone
- **Chips**: Nvidia dominates GPUs, but custom ASICs loom
- **Data Centers:** The core of AI power and scale
- **Opportunity/Risk:** $1 T market vs. overbuild & energy strain

- **Trends:** "Long-thinking" models and agents drive compute demand
- **Sustainability:** Energy, water, carbon now cap scalability
- **Winners/Losers:** Speed + sustainability = survival
- **CoreWeave Now:** Leasing fuels speed but limits control
- **Strategy:** Brownfield retrofits = fast, cheap, independent

Works Cited (A-F)

Amazon Web Services. (2025, September 4). *AWS to show new generative AI innovations at IBC 2025.* https://aws.amazon.com/blogs/media/aws-to-show-new-generative-ai-news-distribution-innovations-at-ibc-2025/
Amazon Web Services. (2025, September 8). *AWS weekly roundup (Sept 8, 2025).* https://aws.amazon.com/blogs/aws/aws-weekly-roundup-aws-transform-amazon-neptune-and-more-september-8-2025/
Amazon Web Services. (2025, October 2). *AWS at IBC 2025: Agentic AI and a decade of AWS Elemental.* https://aws.amazon.com/blogs/media/aws-at-ibc-2025-agentic-ai-and-a-decade-of-aws-elemental/
Amazon Web Services. (2025, October 9). *Serverless ICYMI Q3 2025.* https://aws.amazon.com/blogs/compute/serverless-icymi-q3-2025/
Axios. (2025, October 7). *AI power demand to grow tenfold by 2030, new report says.* https://www.axios.com/2025/10/07/ai-power-cost-demand-future
Bain & Company. (2025, October 24). *AI data center forecast: From scramble to strategy.* https://www.bain.com/insights/ai-data-center-forecast-from-scramble-to-strategy-snap-chart/
Barron's. (2025, October 16). *CoreWeave to work with Nvidia-backed start-up on Texas AI data center project.* https://www.barrons.com/articles/coreweave-nvidia-start-up-texas-data-center-141fb337
Bluefield Research. (2025, July 29). *CoreWeave's recent activity underpins transition for data-center water demand.* https://www.bluefieldresearch.com/research/coreweaves-recent-activity-underpins-transition-for-data-center-water-demand/
Business Insider. (2025, October 27). *America's AI data center dominance, explained in one chart.* https://www.businessinsider.com/data-center-chart-america-dominant-position-torsten-slok-apollo-2025-10
Business Insider. (2025, October 22). *Big Tech's AI capex surge fuels data center wars.* https://www.businessinsider.com/big-tech-ai-capex-infrastructure-data-center-wars-2025-10
CBRE. (2025, August). *North America data center trends H1 2025.* https://www.cbre.com/insights/reports/north-america-data-center-trends-h1-2025
Chain.io. (2024, December 12). *What is vertical integration in the supply chain?* https://chain.io/vertical-integration-supply-chain/
CoreWeave. (2025). *AI infrastructure.* https://www.coreweave.com/ai-infrastructure
CoreWeave. (2025). *GPU compute.* https://www.coreweave.com/products/gpu-compute
CoreWeave. (2025, March 18). *CoreWeave data center operations: Built for AI.* https://www.coreweave.com/blog/coreweave-data-center-operations-built-for-ai
CoreWeave. (2025, May 20). *How we win the AI race: A U.S. infrastructure strategy on our home soil.* https://www.coreweave.com/blog/how-we-win-the-ai-race-a-u-s-infrastructure-strategy-on-our-home-soil
CoreWeave. (2025, June 10). *Our capacity plans for CoreWeave data centers.* https://www.coreweave.com/blog/our-capacity-plans-for-coreweave-data-centers
CoreWeave. (2025, September 9). *CoreWeave launches ventures group to invest in future of AI.* https://www.coreweave.com/news/coreweave-launches-ventures-group-to-invest-in-future-of-ai
CoreWeave. (2025, July 25). *CoreWeave announces two initial data centers in the UK are now operational.* https://www.coreweave.com/news/coreweave-announces-two-initial-data-centers-in-the-uk-are-now-operational
CoreWeave Investors. (2025, August 14). *CoreWeave reports strong second quarter 2025 results.* https://investors.coreweave.com/news/news-details/2025/CoreWeave-Reports-Strong-Second-Quarter-2025-Results/
Data Center Dynamics. (2024, November 21). *The most sustainable data center is the one that's already built: The business case for a retrofit-first mandate.* https://www.datacenterdynamics.com/en/opinions/the-most-sustainable-data-center-is-the-one-thats-already-built-the-business-case-for-a-retrofit-first-mandate/
Data Center Dynamics. (2025, August 1). *CoreWeave signs on for another 150MW of capacity at Applied Digital data center campus.* https://www.datacenterdynamics.com/en/news/coreweave-signs-on-for-another-150mw-of-capacity-at-applied-digital-data-center-campus/
Data Center Dynamics. (2025, September 3). *Global data center power demand set to rise by 50% by 2027— report.* https://www.datacenterdynamics.com/en/news/global-data-center-power-demand-set-to-rise-by-50-by-2027-report/
Data Center Dynamics. (2025, September 10). *CoreWeave ventures launched to back promising AI companies.* https://www.datacenterdynamics.com/en/news/coreweave-ventures-launched-to-back-promising-ai-companies/
Data Center Dynamics. (2025, September 15). *Nvidia to purchase unsold compute capacity from CoreWeave for $6.3bn.* https://www.datacenterdynamics.com/en/news/nvidia-to-purchase-unsold-compute-capacity-from-coreweave-for-63bn/
Data Center Dynamics. (2025, October 1). *Bridging the power gap: Why data centers can't wait for the grid.* https://www.datacenterdynamics.com/en/opinions/bridging-the-power-gap-why-data-centers-cant-wait-for-the-grid/
Data Center Dynamics. (2025, October 15). *S&P Global: US data centers to require 22% more grid-based power by end of 2025.* https://www.datacenterdynamics.com/en/news/sp-global-us-data-centers-to-require-22-more-grid-based-power-by-end-of-2025/
Data Center Dynamics. (2025, October 22). *CoreWeave's CEO says Core Scientific acquisition is "nice to have," not "need to have."* https://www.datacenterdynamics.com/en/news/coreweaves-ceo-says-core-scientific-acquisition-is-nice-to-have-not-a-need-to-have-in-response-to-growing-opposition/
Data Center Dynamics. (2025, October 27). *First 50MW ready for service at Applied Digital Ellendale (CoreWeave 400MW agreement).* https://www.datacenterdynamics.com/en/news/first-50mw-ready-for-service-at-applied-digital-data-center-in-ellendale-north-dakota/
Data Center Dynamics. (2025, November 1). *Let's talk behind-the-meter power for AI data centers.* https://www.datacenterdynamics.com/en/product-news/lets-talk-behind-the-meter-island-power-for-ai-data-centers/
Deutsche Bank Research. (2025, September 29). *AI and energy sectors: More intertwined than ever.* https://www.dbresearch.com/PROD/RI-PROD/PROD0000000000604341/AI_and_energy_sectors__more_intertwined_than_ever.pdf
Deloitte. (2025). *Future-ready AI infrastructure: Accelerating digital transformation.* https://www.deloitte.com/us/en/insights/topics/digital-transformation/future-ready-ai-infrastructure.html
Deloitte. (2025). *Data center infrastructure and artificial intelligence.* https://www.deloitte.com/us/en/insights/industry/power-and-utilities/data-center-infrastructure-artificial-intelligence.html
Energy Connects. (2025, April 10). *IEA: AI to accelerate energy innovation and drive surging electricity demand.* https://www.energyconnects.com/opinion/features/2025/april/iea-ai-to-accelerate-energy-innovation-and-drive-surging-electricity-demand/
Epoch. (2025, November 7). *Frontier AI data centers: Open database and methodology.* https://epoch.ai/data/data-centers
Finviz/Zacks Equity Research. (2025, August 6). *How will CoreWeave's data center push power AI innovation?* https://finviz.com/news/207235/how-will-coreweaves-data-center-push-power-ai-innovation
Flaningam, E. (2024, October 13). *A primer on AI data centers.* Generative Value. https://www.generativevalue.com/p/a-primer-on-ai-data-centers
Forbes. (2024, December 11). *The AI chip race: Who can compete with Nvidia?* https://www.forbes.com/sites/emilsayegh/2024/12/11/the-ai-chip-race-who-can-compete-with-nvidia/
FT (Financial Times). (2025). *AI data center map.* https://ig.ft.com/ai-data-centres/
FT (Financial Times). (2025). *Big Tech's data center spending race heats up.* https://www.ft.com/content/be8d4c0c-66ff-4dfd-9b43-af6c0b290ada
FT (Financial Times). (2025). *Inside the global AI buildout.* https://www.ft.com/content/8452bf94-9a41-4040-913f-ef1a462d6ea6
FT (Financial Times). (2025). *AI's infrastructure crunch: Power, chips, and the future of compute.* https://www.ft.com/content/efe1e350-62c6-4aa0-a833-f6da01265473

Works Cited (F-Z)

FT (Financial Times). (2025). *CoreWeave's quiet rise in the AI arms race.* https://www.ft.com/content/41bfacb8-4d1e-4f25-bc60-75bf557f1f21

Gartner. (2025). *AI infrastructure market guide (Report ID 4742298).* https://www.gartner.com/document/code/4742298

GlobeNewswire/SNS Insider. (2025, October 31). *Data center market to reach USD 712B by 2033.* https://www.globenewswire.com/news-release/2025/10/31/3178144/0/en/Data-Center-Market-to-Reach-USD-712-00-Billion-by-2033-Driven-by-Cloud-Expansion-and-AI-Powered-Infrastructure-SNS-Insider.html

Harvard Business Review. (2025, June). *The new cloud oligopoly: AI compute providers and strategic risk.* https://hbr.org/2025/06/the-new-cloud-oligopoly-ai-compute-providers-and-strategic-risk

IDC. (2025). *U.S. data center outlook and AI infrastructure report (Doc #US51544025).* https://www.idc.com/getdoc.jsp?containerId=US51544025

International Energy Agency. (2025). *Energy and AI.* https://iea.blob.core.windows.net/assets/601eaec9-ba91-4623-819b-4ded331ec9e8/EnergyandAI.pdf

International Energy Agency. (2025). *Energy demand from AI.* https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai

McKinsey & Company. (2025, June 26). *The cost of compute: A $7 trillion race to scale data centers.* https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-cost-of-compute-a-7-trillion-dollar-race-to-scale-data-centers

McKinsey & Company. (2025, October). *Superagency in the workplace: Empowering people to unlock AI's full potential.* https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/superagency-in-the-workplace-empowering-people-to-unlock-ais-full-potential-at-work

McKinsey & Company. (2025). *The data center balance: How U.S. states can navigate opportunities and challenges.* https://www.mckinsey.com/industries/public-sector/our-insights/the-data-center-balance-how-us-states-can-navigate-the-opportunities-and-challenges

MIT Sloan Management Review. (2025). *Vertical integration in the AI infrastructure stack.* https://sloanreview.mit.edu/article/vertical-integration-in-the-ai-infrastructure-stack/

Nasdaq. (2025, October 28). *How will CoreWeave's data center push power AI innovation?* https://www.nasdaq.com/articles/how-will-coreweaves-data-center-push-power-ai-innovation

New York Times. (2025, October 31). *AI spending is accelerating.* https://www.nytimes.com/2025/10/31/technology/ai-spending-accelerating.html

OECD. (2025). *Compute concentration and policy implications (AI Policy Brief).* https://oecd.ai/en/dashboards/policy-briefs/compute-concentration-2025

Pew Research Center. (2025, October 24). *What we know about energy use at U.S. data centers amid the AI boom.* https://www.pewresearch.org/short-reads/2025/10/24/what-we-know-about-energy-use-at-us-data-centers-amid-the-ai-boom/

Precedence Research. (2025). *Artificial intelligence market size, trends, and forecasts.* https://www.precedenxceresearch.com/artificial-intelligence-market

ProQuest. (2025). *AI investment and infrastructure growth reports.* https://www.proquest.com/docview/3147525788/356BBED3C70D4130PQ/

ProQuest. (2025). *Big Tech's data center expansion: Risks and rewards.* https://www.proquest.com/docview/3178311516/356BBED3C70D4130PQ/

Reuters. (2025, September 15). *CoreWeave, Nvidia sign $6.3B cloud capacity backstop.* https://www.reuters.com/business/coreweave-nvidia-sign-63-billion-cloud-computing-capacity-order-2025-09-15/

Reuters. (2025, October 15). *Meta commits $1.5B for AI data center in Texas.* https://www.reuters.com/business/meta-commits-15-billion-ai-data-center-texas-2025-10-15/

Reuters. (2025, October 31). *The great AI buildout shows no sign of slowing.* https://www.reuters.com/legal/transactional/great-ai-buildout-shows-no-sign-slowing-2025-10-31/

SEC. (2025, September 9). *CoreWeave Form 8-K (CRWV-20250909).* https://www.sec.gov/Archives/edgar/data/1769628/000176962825000047/crwv-20250909.htm

SEC. (2025, September 9). *CoreWeave S-1 filing.* https://www.sec.gov/ix?doc=/Archives/edgar/data/1769628/000176962825000032/coreweave-s1.htm

SEC. (2025, January 28). *Nvidia Form 10-K.* https://www.sec.gov/ix?doc=/Archives/edgar/data/1045810/000104581025000012/nvda-20250128.htm

Statista. (2025). *Worldwide market share of leading cloud infrastructure service providers.* https://www.statista.com/chart/18819/worldwide-market-share-of-leading-cloud-infrastructure-service-providers/

Tom's Hardware. (2025, July 8). *CoreWeave acquires Core Scientific in long-awaited move.* https://www.tomshardware.com/tech-industry/artificial-intelligence/ai-hyperscaler-buys-its-cryptomining-partner-for-its-ai-gpus-and-data-center-infrastructure-coreweave-acquires-core-scientific-in-long-awaited-move

Troutman Pepper. (2025, July 23). *Navigating contractual considerations in the AI data center construction boom.* https://www.troutman.com/insights/navigating-contractual-considerations-in-the-ai-data-center-construction-boom/

U.S. Environmental Protection Agency. (2025). *About brownfields.* https://www.epa.gov/brownfields/about