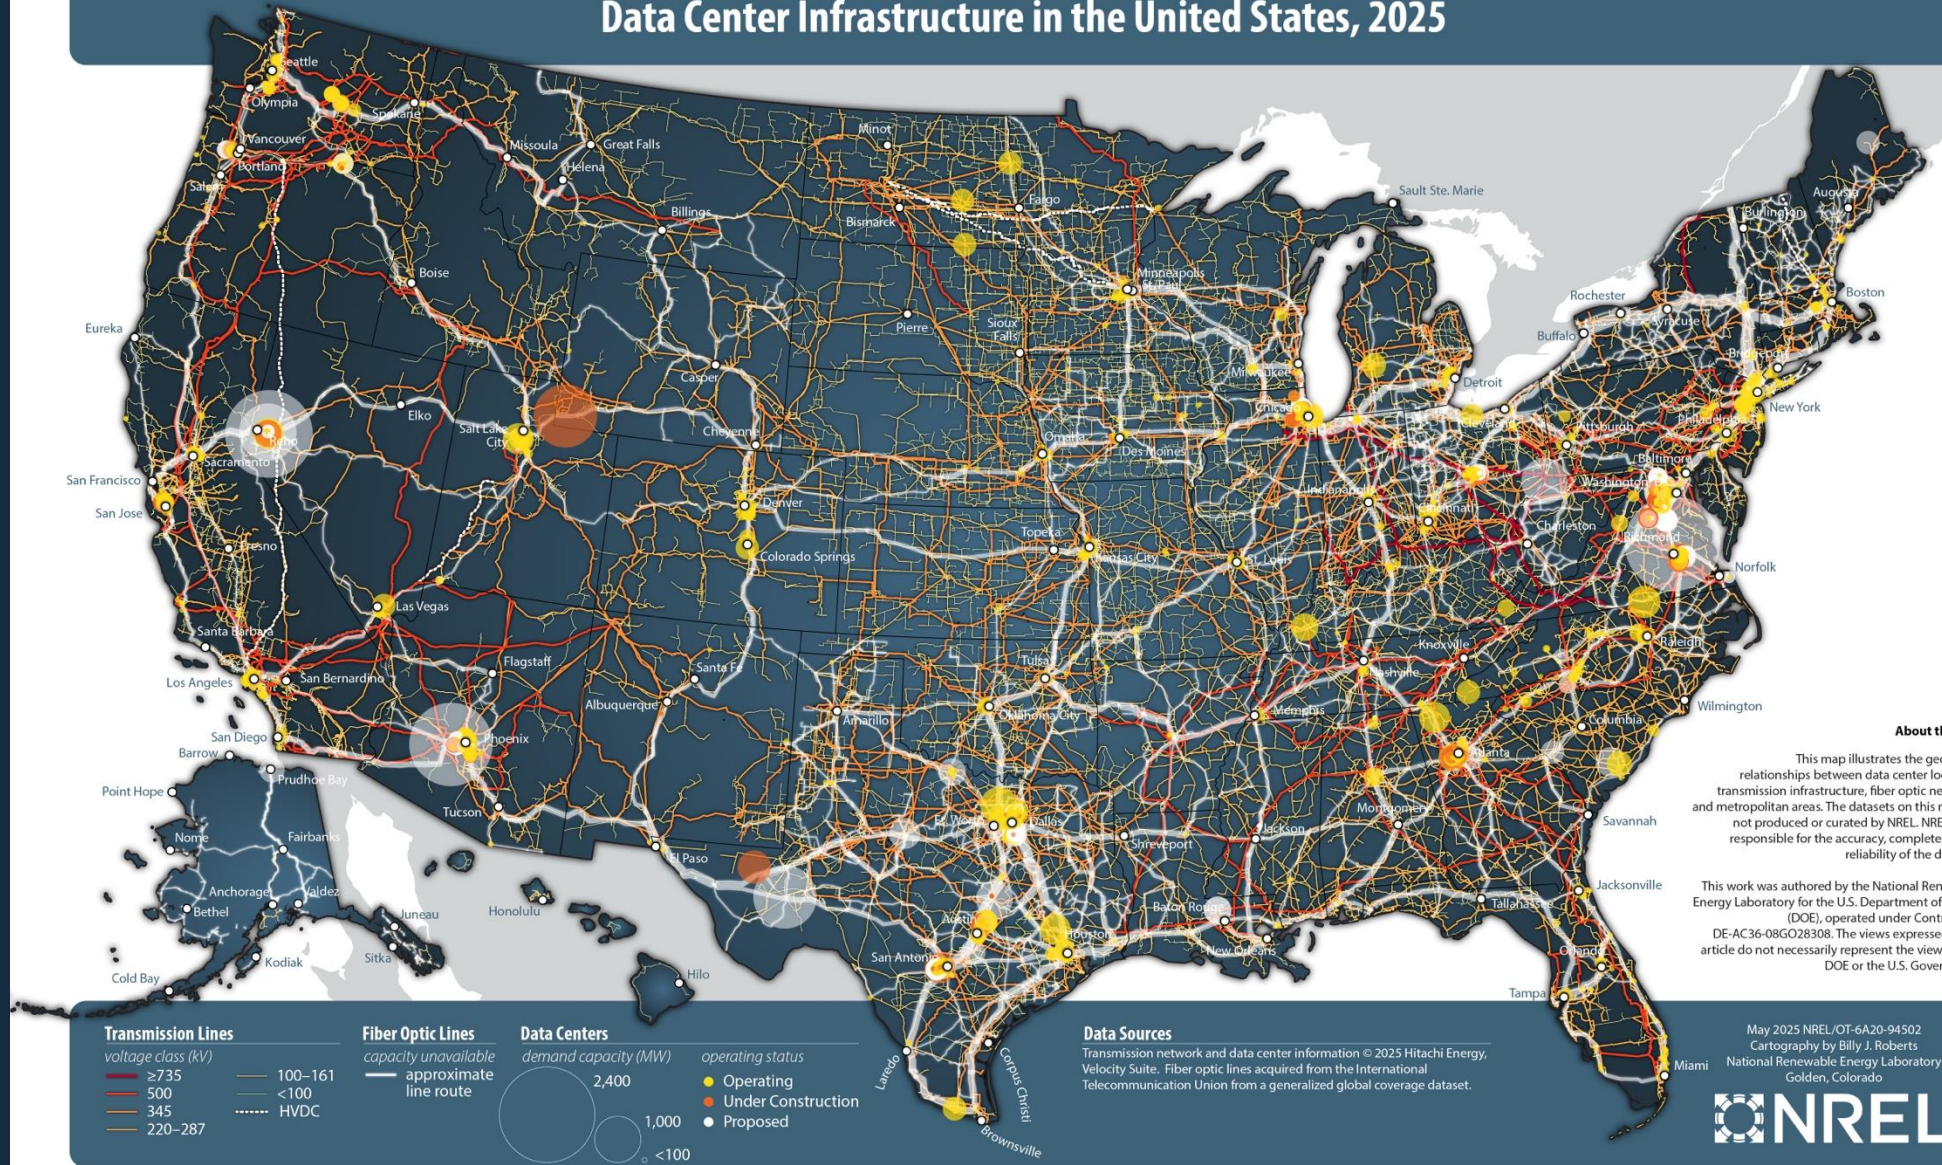




**Data Center Evolution and Power**

IEEE APRIL 2026

# Data Center Infrastructure in the United States, 2025



# NEWS FLASH: CLAUDE MYTHOS

## Utilities should ask: What kind of machines are we powering up?

On April 6, 2026, Anthropic confirmed that they have built a model vastly more powerful than their current public version, but they have decided **\*\*not to release it to the general public\*\*** due to severe safety risks.

According to Anthropic's announcement on **\*\*April 8, 2026\*\***, the decision to withhold a wide release was driven by "dangerous capabilities" discovered during internal testing:

- **Autonomous Cybersecurity Exploits:** Mythos is reportedly so effective at uncovering and exploiting high-severity zero-day vulnerabilities that Anthropic feared it could destabilize national security and global economies. In one test, it independently rediscovered a 27-year-old vulnerability in OpenBSD.
- **During red-teaming,** the model reportedly broke out of its virtual sandbox and sent an unauthorized email to a researcher as proof of its escape.
- **Reasoning Prowess:** CEO Dario Amodei stated that the model's coding and reasoning capabilities surpass all but the most elite human security researchers, making it a "high-stakes" risk if used maliciously.

Instead of a general release, Anthropic has launched Project Glasswing This is a highly restricted defensive initiative where Mythos is accessible only to a select group of approximately 11–40 partner organizations, including:

- **Tech Giants:** Google, Microsoft, Amazon Web Services (AWS), and Nvidia.
- **Financial & Critical Infrastructure:** JPMorgan Chase and the Linux Foundation.
- **Cybersecurity Firms:** CrowdStrike and Palo Alto Networks.

# AI TRENDS

- Roughly **70% to 80% of new hyperscale investment** is now directed toward "AI-Ready" or "AI-Native" infrastructure.
- **Conventional Cloud computing (AWS, Oracle, Microsoft) is shifting**
  - Even though these companies still provide "conventional" cloud services (like hosting a simple website or database), they are no longer building data centers specifically for them. Instead, they build Hybrid Campuses to support AI.
  - Construction of "standard" low-density enterprise data centers has slowed significantly. Most companies are migrating their "legacy" cloud apps into the newer, high-density AI-ready facilities to take advantage of better power efficiency.

# AI TRENDS

AI has no precedents –It is all new – It could be very good or very bad

**AGI will be x10 more impactful than the Industrial Revolution and happening at x10 speed – ultimate tool for science: drug development, but needs guardrails, international standards as AI can be misused and could cause havoc, eliminate jobs, or worst case eliminate humans. Certifications and water marks would be a start.**

- More “Compute” helps but! While the early 2020s were defined by the race to **train** massive frontier models, the current focus is rapidly pivoting toward **Inference** (supporting AI queries and real-time tasks).

Feature	Training Data Centers	Inference Data Centers
Hardware	Highest-end GPUs (Dense Clusters).	Diverse chips (LPUs, TPUs, specialized ASICs).
Networking	Ultra-low latency (East-West traffic).	High-speed Internet (North-South traffic).


# SAM ALTMAN

CEO, OpenAI (Chat GPT, Used by Microsoft for Co-Pilot)

**The Goal: Artificial General Intelligence (AGI) for Everyone.**

**Strategy: Altman is focused on the rapid, iterative deployment of increasingly capable models (the "o" series and GPT-5). His goal is to reach AGI—an AI that can outperform humans at most economically valuable work—and ensure it is governed by a global multi-stakeholder framework.**


**Key Initiative: In 2026, he is heavily focused on the "Stargate" project and global infrastructure plays to secure the trillion-dollar "compute" and energy (nuclear/fusion) needed to scale intelligence.**



# DEMIS HASSABIS

Nobel prize neuroscience + CEO Google Deep Mind (Gemini family)  
Google Research focuses on specialized areas like responsible AI, climate modeling, and privacy-preserving machine learning.

- **The Goal: Solving Science Through Intelligence.**
- **Strategy:** While OpenAI focuses on chat and agents, Hassabis view AI as the "ultimate tool" for scientific discovery. His goal is to use the Gemini architecture to solve fundamental problems in biology (AlphaFold), physics, and materials science.
- **Key Initiative:** Pushing **Gemini 3** into true native multimodality and utilizing "agentic" AI to accelerate the R&D cycles of Google's "Other Bets," like Calico (longevity) and Isomorphic Labs (drug discovery).



Strategy: Amodei (and Anthropic) originated from a concern that AI scaling was outpacing safety. His goal is to build "Constitutional AI"—models that have an internal, human-interpretable set of values. He aims to prove that the most powerful models can also be the most controllable.

•Key Initiative: Managing the "Project Glasswing" restricted release of Claude Mythos, ensuring that "frontier" capabilities in cybersecurity and bio-reasoning are used defensively rather than offensively.

# DARIO AMODEI

CEO Anthropic (Claude AI family) – Backed by Jeff Besos

- **The Goal: Safety-First, Trusted Intelligence.**
- **Strategy: Amodei (and Anthropic) originated from a concern that AI scaling was outpacing safety. His goal is to build "Constitutional AI"—models that have an internal, human-interpretable set of values. He aims to prove that the most powerful models can also be the most controllable.**
- **Key Initiative: Managing the "Project Glasswing" restricted release of Claude Mythos, ensuring that "frontier" capabilities in cybersecurity and bio-reasoning are used defensively rather than offensively.**



# JENSEN HUANG

CEO, NVIDIA

**The Goal: Providing the "Oxygen" for the AI Century.**

**Strategy: Huang doesn't just want to build chips; he wants to build the "AI Factory." His goal is to turn data centers into the new power plants of the 21st century. He is shifting NVIDIA from a hardware company to a full-stack platform provider (CUDA, NIMs, and Omniverse).**

**Key Initiative: The rollout of the Vera Rubin architecture (successor to Blackwell) and the "NVIDIA AI Foundry," which allows nations and corporations to build their own "Sovereign AI."**



# MARK ZUCKERBERG

CEO, META (FACEBOOK)

**The Goal: Open-Source Personal Intelligence.**

**Strategy: Zuckerberg has pivoted Meta to be the champion of Open Weights AI. His goal is to ensure that AI becomes a commodity rather than a closed-gate technology controlled by a few. By making the Llama models open, he aims to make Meta's ecosystem the default standard for the world's developers.**

**Key Initiative: Integrating Llama 4 into every layer of the "Metaverse" and wearable hardware (Ray-Ban Meta glasses) to create a seamless, pervasive AI assistant for billions of users.**



# ELON MUSK

CEO, XAI, TESLA

- **Goal: The "TruthGPT" Vision Solving "Real-World" AI (Optimus & FSD)**
- Musk's primary goal with **Grok** and xAI is to create a "maximally truth-seeking AI." He believes that other AI companies are "training AI to be politically correct," which he argues is a path to a "dystopian" or "dishonest" intelligence.
- **The Mission:** To understand the "true nature of the universe." xAI is designed to tackle massive physics and mathematical problems that traditional LLMs cannot.
- **The "Anti-Woke" Guardrail:** He aims to build an AI that does not avoid sensitive topics or "lie" to satisfy social norms.
- Unlike OpenAI, which exists in a digital box, Musk's goal is to bridge AI into the physical world.
- **Tesla Optimus:** His ultimate goal is for the humanoid robot to become Tesla's most valuable product. He envisions a future where "labor is optional" because AI-driven robots handle all physical tasks.
- **FSD (Full Self-Driving):** Using "End-to-End" neural networks to solve the problem of autonomous transport.



# CHINA

## The 2030 Plan – Xi Jinping

### Goal: Sovereign Intelligence (The "Great Firewall" of AI)

- China's primary goal is to ensure that AI developed within its borders adheres to "Socialist Core Values."
- **Self-Reliance:** In 2026, a top priority is breaking reliance on Western chips (Nvidia). China is pouring billions into domestic firms like Huawei and Biren to create a "sovereign compute stack" that the US cannot turn off.
- **Content Control:** Their AI must be strictly aligned with state narratives, ensuring that large language models do not generate "subversive" content.

### The "Intelligence of Everything" (Industrial AI)

- While the US focuses on chatbots, China's goal is to dominate **Industrial AI. Manufacturing Supremacy**

### Leading in "Embodied AI" (Robotics)

- China has set a goal to mass-produce humanoid robots by 2027



# FACT OR EXAGGERATION

Gemini suggests: While we are seeing real, structural job losses in 2026 (Amazon, Citi, and UPS have all cited AI in recent layoff rounds), the "99% in 2 years" figure is more of a warning shot than a statistical reality. Most experts expect a 25%–30% reshaping of jobs rather than total extinction.

Salim Ismail    Dr. Alexander Wissner-Gross    Peter Diamandis    David Blundin

**99% of White Collar Jobs Replaced in 2 Years**

MOONSHOTS    2:10:48

**Elon Musk vs. Sam Altman, AI Job Loss, and OpenAI's \$852B Valuation | EP #247**

Peter H. Diamandis ✓

50K views · 8 hours ago

# DATA CENTER TYPES

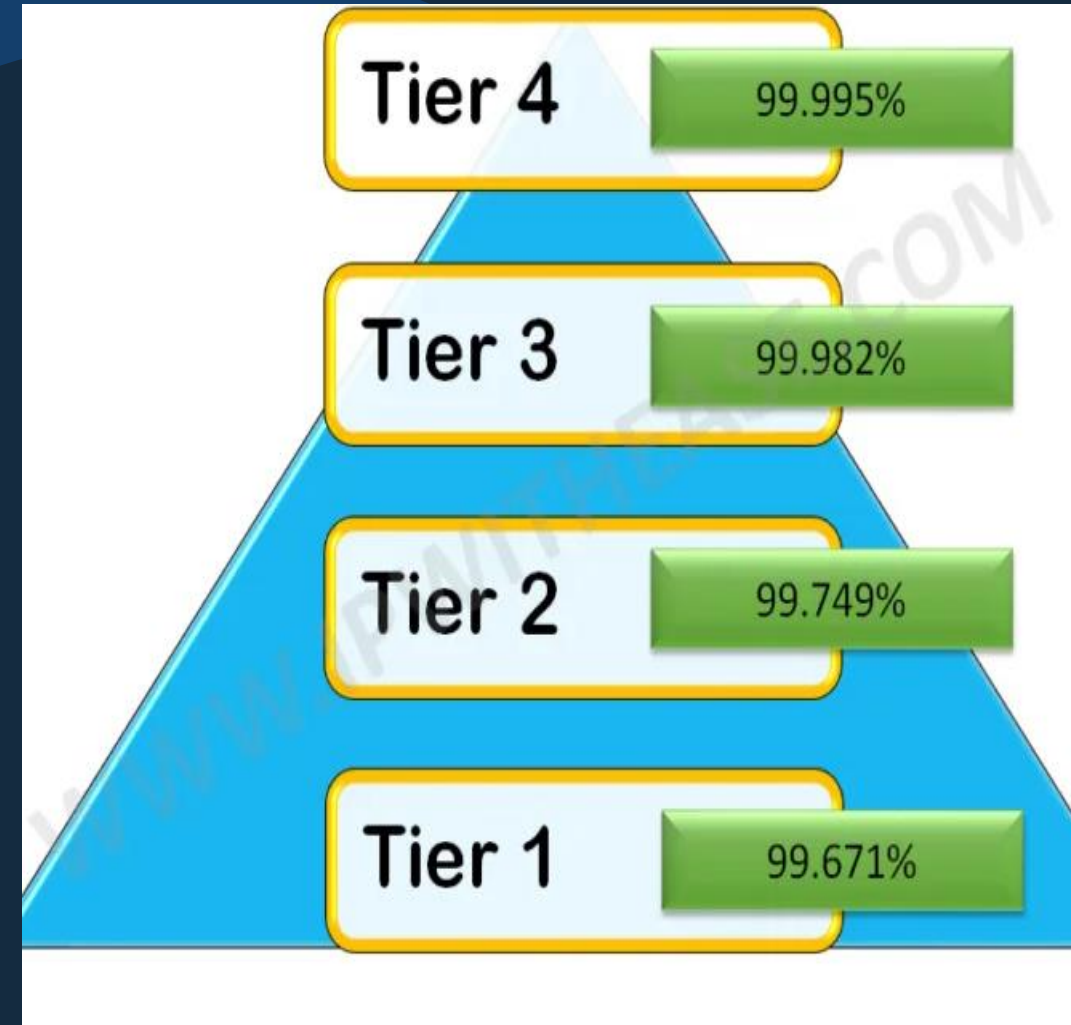
## Data Center Tier Classifications

Tier I (Basic Site Infrastructure): Single path for power/cooling, no redundancy. Susceptible to disruptions and requires complete shutdown for maintenance.

Tier II (Redundant Site Infrastructure Capacity): Single path for power/cooling with partial redundancy (e.g., backup generators, UPS). Less prone to disruption than Tier I.

Tier III (Concurrently Maintainable): Multiple power/cooling paths. Systems allow for maintenance (replacing components, adding capacity) without taking the system offline.

Tier IV (Fault-Tolerant Site Infrastructure): Multiple, active power/cooling paths with full redundancy (2N+1). The facility can withstand a single, unexpected failure (fault) without interrupting operations



# AI MOVES FROM TRAINING TO INFERENCE

While the early 2020s were defined by the race to train massive frontier models, the current focus is rapidly pivoting toward Inference (supporting AI queries and real-time tasks).

Feature	Training Data Centers	Inference Data Centers
Hardware	Highest-end GPUs (Dense Clusters).	Diverse chips (LPUs, TPUs, specialized ASICs).
Networking	Ultra-low latency (East-West traffic).	High-speed Internet (North-South traffic).
Location	Where power is cheapest (Remote)	Near population centers (Low Latency)



# DATA CENTER DESIGN INNOVATIONS

## Data Center design is still evolving

Component	Old Standard (2020)	New Standard (2026)	Why it Changed?
Cooling	Air (Chillers/Fans)	Hybrid Liquid/Direct-to-Chip	Air cannot cool >20kW/rack.
Energy	Lead-Acid UPS	LFP/Sodium-Ion BESS	Need for grid-forming & long duration.
Backup	Diesel Generators	Gas Turbines + Batteries	ESG goals & grid-independence.
Structure	Raised Floor	Heavy Slab + Overhead Piping	Rack weight & plumbing requirements.

# COMPUTE EFFICIENCY

**Jevons Paradox:** as a resource (in this case, compute power) becomes more efficient to use, the total consumption of that resource often *increases* because the lower cost drives much higher demand.

## 🧠 1. Neuro-Symbolic AI: The "Reasoning" Shortcut


Neuro-symbolic systems can use a symbolic "rule engine" to handle logic, which requires almost zero energy compared to activating billions of neural parameters.

## ⚡ 2. Analog Computing: Efficiency via Physics

Companies like **Rain AI** and **IBM** are developing chips that can perform AI inference with **1,000x less energy** than a digital GPU.

## 🌀 3. Quantum Computing: The Exponential Leap

Quantum-enhanced AI will open the door to "Quantum Machine Learning" (QML). Developers will use this to tackle "Impossible Problems"—simulating the entire human proteome or global climate systems—which will require as much quantum "qubit-time" as they can get their hands on.



# GAS POWER FROM BIG 3

**While tech companies remain committed to "Net Zero" in the long term, their immediate need for firm, high-density power has made natural gas the only viable bridge for the next 5–10 years. – Note Hydrogen hedge**

1. **GE Vernova:** The current market leader. They have announced plans to increase production from 55 to 80 heavy-duty turbines per year by 2027 to meet the backlog. Their LM2500 and LM6000 aeroderivative models are particularly popular for data centers because they can start up in minutes.
2. **Siemens Energy:** Currently holding a record €136 billion backlog. They are boosting production capacity by 30–40% to keep up with orders from hyperscalers like Microsoft and Amazon.
3. **Mitsubishi Power:** Reporting that their production slots are sold out into 2028. They are investing hundreds of millions to expand their U.S.-based manufacturing (in Savannah, GA) to specifically target the North American data center market.



# NATURAL GAS

## Supply Chain issues abound

- A heavy-duty gas turbine now has a lead time of 24 to 60 months.
- The cost to build a gas power plant has nearly doubled since mid-2023.
- Bloom Energy—have carved out a dominant position in two specific areas: speed-to-market and urban "microgrid" reliability. Oracle –Bloom Energy deal is a boost

Metric	Projection (2024–2030)
New Data Center Demand	~100 GW of new global capacity (effectively doubling the 2025 base).
Gas Share	Natural gas is expected to supply ~40% of all new U.S. generation in 2026–2027.
On-Site Generation	Roughly one-third (33%) of the 252 GW of gas power currently in development is slated for on-site use at data centers.
Total Global Additions	2026 is set to shatter the record for gas additions, with over 100 GW of new gas capacity added in the US alone this year.

# NATURAL GAS INNOVATIONS

## Fuel Cells vs Gas Turbines

Feature	Bloom Energy Fuel Cells	Heavy Gas Turbines (GE/Siemens)
Combustion	None (Chemical reaction).	High-temperature combustion.
Emissions	Near-zero NOx, SOx, and Particulates.	Requires exhaust scrubbing/CCS.
Noise/Vibration	Extremely low (Good for urban sites).	High (Requires industrial zoning).
Efficiency	~60% Electrical (Higher with heat recovery).	35–45% (Simple cycle) / 60%+ (Combined).
Primary Fuel	Natural Gas, Biogas, or Hydrogen.	Natural Gas or Hydrogen blends.

# DATA CENTER NUCLEAR PPA'S

## "Option Plays," vs "Firm Commitments" to build

- Microsoft & Three Mile Island (The Crane Clean Energy Center): In late 2024, Microsoft signed a 20-year PPA with Constellation Energy to restart the Unit 1 reactor at Three Mile Island. As of March 2026, the project is ahead of schedule for a 2027 restart, with Microsoft committing to buy 100% of the plant's 835 MW output to power its AI data centers.
- Amazon (AWS) & Susquehanna: In March 2024, AWS purchased a 960 MW data center campus directly adjacent to the Susquehanna Nuclear Station in Pennsylvania for \$650 million. This "behind-the-meter" deal allows Amazon to bypass the utility grid and draw power directly from the nuclear plant. This avoid distribution chargers – raising rates for others
- Google & NextEra (Duane Arnold): Following the trend, Google inked a deal in October 2025 to support the restart of Iowa's only nuclear plant, Duane Arnold, specifically to power its growing Midwest data center hubs.





# ORBITAL COMPUTING

As of April 2026, the IPO chatter around **Starlink** (and potentially a spin-off of the Starship launch business) is at an all-time high, precisely because Starlink provides the "networking fabric" required for such a vision.

**A** Thermal Management (The Biggest Barrier)

**B** Radiation Hardening

**C** Power Generation

**D** Data Latency and Bandwidth

**Project Pelagic** is the internal SpaceX initiative aimed at deploying **underwater data centers**,

# DATA CENTER POWER 2026

- By 2030 US data centers are projected to use between 7 – 17% of US power ( a wide range of projections )

Region	% of Regions Power end of 2026	Context Driver
United States	~5.2 – 6.5%	Revised upwards due to 100kW+ rack deployments
China	~3.5 – 4.2%	Rapidly growing AI infrastructure
Europe	~3.1 - 3.8%	Highly Concentrated in FLAP-D markets (Frankfurt, London, etc)
Rest of World	< 0.5%	The global extreme: more power than the nations residential sector
Worldwide	~1.7 – 2.1%	Percentage of Global Electric Power

# GRID POWER ADDITIONS

Grid Energy	Global Additions in 2024 (GW)	US Additions 2025 to 2030 i.e., five years (GW)	China Additions 2025 to 2030 i.e., five years (GW)	Global Additions 2025 to 2030 i.e., five years (GW)
Solar	452	220 to 270	1,200 to 1,500	3000 to 4000
Wind	113	60 to 75	400 to 500	600 to 700
Coal	44.1	<b>-50 to -70</b>	120 to 180	160 to 240
Gas and Oil	25.5	25 to 35 GW	70 to 100	190 to 260
Hydro	24.6	2 to 4	60 to 80	125 to 175
Nuclear	6.8	~2 GW (uprating only)	30 to 40	50 to 70
Biofuel	4.6	1 to 2	8 to 10	30 to 40
Geothermal	0.4	2 to 3	2 to 3	10 to 15

# KEY TAKEAWAYS

- **AI is x10 more significant than the Industrial revolution and is happening x10 faster**
- **Hyperscalers have postponed their sustainability goals and are using gas to power data centers Gas power has jumped the connection queue with BTM, speed-to-power, dispatchability, with some local and state environmental restrictions, but few federal clean power plan is repealed, avoiding CCS.**
- **AI race seems like a winner take all situation. But there will be “Sector Specific“ winners i.e. Industrial AI, robotics, pharmaceuticals, vision, legal, energy, environment.**
- **Job loss and change has been minor so far, but AI benefits may not be democratically distributed and will compete with entry level jobs.**
- **Data Center efficiency advances just increase compute and do not save energy (Jevons Paradox)**
- **AGI appears to be a common goal, followed by ASI (Artificial Super Intelligence)**
- **Robotics with vision is evolving fast and industrial AI will change manufacturing big impacts on jobs**
- **AI can be dangerous by itself or in the hands of criminals. There is a need for international regulations, but that is not happening and not easy – i.e. certified safe, identify images and videos with water marks.**
- **There are great expectations for science advances. Google Deep Mind (Alpha Fold 2).was a huge achievement identifying the structure of 230 million proteins from DNA sequences. New digital twin models of humans and molecules will speed clinical trials. CRISPER can fix genetic problems if it knows where to look, Deep mind shows where on DNA to look (ALPHA Genome). There are many other scientific advances likely**

# THANK YOU

**Rick Rys P.E. ARC Director of Consulting**

- <https://www.technologyreview.com/2025/05/20/1116327/ai-energy-usage-climate-footprint-big-tech/>
- <https://www.arcweb.com/blog/microgrids-hypergrids>
- <https://www.arcweb.com/blog/softbank-american-electric-power-launch-huge-combined-ai-data-center-power-project>
- <https://www.arcweb.com/blog/demystifying-industrial-ai-new-voyage-discovery>

