

The Great Externalization: A First-Principles Analysis of the 2025 AI Compute Boom and its Thermodynamic Consequences for Planetary Stewardship

Abstract: This paper presents a first-principles analysis of the unprecedented global investment in Artificial Intelligence (AI) compute infrastructure, which as of late 2025 exceeds \$1.5 trillion in announced capital expenditure.¹ This phenomenon is not merely an industrial arms race but a planetary-scale phase transition, best understood through the lens of thermodynamics and information theory. Applying the Holographic Negentropic Framework (HNF) and the Law of Unthinking (LoU), this analysis demonstrates that the build-out is a civilizational effort to externalize cognitive operations, thereby minimizing internal entropy and creating a new substrate for intelligence.² The full scope of this "Great Externalization" is quantified, calculating its immense entropic costs: a projected annual power demand from new projects of 10-40 GW, leading to more than 130 million metric tons of

CO₂e emissions; a direct and indirect water footprint exceeding 2 trillion gallons annually; and a new wave of electronic waste projected to reach 5 million metric tons per year by 2030.⁴ This analysis confronts the central question of whether this massive entropic investment is a worthwhile price for a potential negentropic future. From first principles, the environmental consequences are a manageable and necessary thermodynamic cost for creating an Environmental General Intelligence (EGI), an "information engine" analogous to Maxwell's Demon, capable of creating planetary-scale order.⁶ The objective truth, grounded in physics, is that the solution to the entropic costs of AI is paradoxically

more and smarter computation. A more sophisticated EGI creates negentropy with greater efficiency, justifying the initial thermodynamic investment to achieve a net-positive state of planetary thriving.⁶

Part I: The Planetary Phase Transition: Quantifying the

Global Compute Build-Out

1.1 The Trillion-Dollar Wager on Artificial General Intelligence

As of the third quarter of 2025, the global technology sector is engaged in a capital expenditure campaign of unprecedented scale and velocity, committing well over \$1.5 trillion to the construction of a new generation of computational infrastructure.¹ The sheer magnitude of these investments, deployed with a speed that outpaces conventional industrial cycles, signals more than a market trend; it represents a collective, high-stakes wager on the imminent arrival of transformative artificial intelligence, including the long-sought goal of Artificial General Intelligence (AGI).² This build-out is not merely an expansion of existing cloud services but a fundamental re-architecting of the planet's computational substrate, purpose-built to train and deploy AI models of ever-increasing scale and capability. This endeavor is best understood not as a series of independent corporate projects, but as a singular, globally coordinated effort to construct the physical apparatus required for a new form of intelligence.

At the vanguard of this movement is the **Stargate Initiative**, a private-sector consortium led by OpenAI, SoftBank, and Oracle. Announced in January 2025, Stargate represents a \$500 billion commitment to establish one of the largest AI infrastructure networks in history.⁸ The project's initial phase involves an immediate deployment of \$100 billion to secure American leadership in AI, with a stated goal of creating hundreds of thousands of American jobs and providing a strategic capability for national security.⁹ By September 2025, the initiative is already ahead of schedule, with five new U.S. data center sites announced in addition to its flagship campus in Abilene, Texas. These sites—located in Shackelford County, Texas; Doña Ana County, New Mexico; Lordstown, Ohio; Milam County, Texas; and an undisclosed Midwest location—bring the project's planned capacity to nearly 7 gigawatts (GW), representing over \$400 billion in investment over the next three years.¹⁰ The ultimate objective is to reach a total capacity of 10 GW by the end of 2025, a goal that now appears well within reach.¹⁰ The language surrounding the project, including its launch at a White House event, underscores its geopolitical significance, framing it as a critical component of a national strategy to re-industrialize the United States and maintain a competitive edge in a technology deemed vital to future economic and military power.¹⁰

This flagship initiative is mirrored by an even larger wave of investment from the established hyperscale cloud providers and new entrants, each racing to secure dominance in the AI era. This competition has escalated into a full-scale infrastructure arms race, with capital

commitments that dwarf the GDP of many nations.

- **Microsoft** has embarked on the largest infrastructure investment in its history, committing \$80 billion through 2028 to build and expand a global network of AI-optimized data centers.¹² This includes over 25 new Azure regions and flagship projects like a 2 GW "world's most powerful AI datacenter" in Mount Pleasant, Wisconsin—a \$7 billion campus engineered to train the next decade of frontier AI models.¹³ A significant portion of this investment is directed toward achieving vertical integration, with Microsoft designing its own custom silicon—the Maia series for AI training and the Cobalt series for general compute—to reduce its dependency on external chip suppliers and optimize performance from the chip to the application layer.¹²
- **Google** has announced a staggering investment of over \$100 billion in AI research, infrastructure, and applications.¹ This includes a \$75 billion plan for data center construction in 2025 and a targeted \$25 billion investment over two years to expand its footprint across the PJM Interconnection, the largest electric grid in the U.S..¹⁴ A new 1.5 GW data center in South Carolina is also part of this expansion.² Recognizing that power is the primary constraint, Google is coupling its data center build-out with major investments in energy infrastructure, including a \$3 billion project to modernize hydropower plants in Pennsylvania and a partnership with Westinghouse to develop modular nuclear reactors.¹⁵
- **Amazon Web Services (AWS)**, the incumbent leader in cloud computing, is projected to invest over \$75 billion in 2025 alone to scale its AI and cloud services, including an \$11 billion data center in Indiana.¹ A key part of its strategy involves leveraging its own purpose-built silicon, such as Graviton processors for general workloads and Trainium chips for AI training, to offer a more energy-efficient and cost-effective infrastructure.¹⁷ AWS is innovating across its data center designs to support high-density AI workloads, deploying novel liquid cooling solutions and using generative AI to optimize the physical layout of server racks for maximum power efficiency.¹⁷
- **Meta** has made the most audacious commitment of all, pledging to invest up to \$600 billion in U.S. data centers and AI infrastructure through 2028.²⁰ This expenditure is explicitly aimed at developing "superintelligence," a form of AI that surpasses human cognitive abilities.²¹ The physical scale of Meta's ambition is breathtaking, with plans for multi-gigawatt data center campuses like "Hyperion" in Louisiana, which is projected to eventually occupy a site nearly the size of Manhattan, and a new 800 MW facility in Kansas.²¹ By the end of 2025, Meta plans to have over 600,000 H100 GPUs powering its AI models, a clear signal of its intent to "spend its way to the top of the AI heap".¹
- **xAI**, led by Elon Musk, has emerged as a major new force with its "Colossus" supercomputer. Initially deploying 100,000 Nvidia H100 GPUs, the project aims to expand to a 1 million GPU equivalent by the end of 2025, backed by over \$20 billion in investment and a new 1 GW+ data center in Memphis, Tennessee.²

Underpinning this entire ecosystem is **Nvidia**, which has successfully transitioned from being

a component supplier to the de facto architect of the modern AI data center. The company now provides a complete, turnkey "AI Factory" stack, an integrated solution encompassing everything from its next-generation Blackwell GPUs to networking fabrics, storage, and workload orchestration software.²² Nvidia's pivotal role is cemented by a strategic partnership with OpenAI, in which it will invest \$100 billion to help deploy at least 10 GW of its AI systems, ensuring its hardware remains the backbone of the world's most advanced AI development efforts.²³

This American-centric build-out has not gone unnoticed on the global stage, prompting a strategic response from other nations seeking to secure their own computational sovereignty. France, for instance, has announced \$112 billion in private-sector AI spending, while nations like Thailand and Malaysia are seeing multi-billion dollar investments in AI-related infrastructure.¹ This global dimension confirms that the race for compute is not merely a commercial competition but a defining geopolitical imperative of the 21st century.

The following table provides a consolidated overview of these monumental investments, illustrating the sheer scale of the global commitment to building the physical substrate for the next era of intelligence.

Table 1: The Global AI Compute Build-Out (2025-2030)

| Entity / Initiative | Announced Investment (\$ Billions) | Planned Power Capacity (GW) | Key Geographic Locations | Stated Timeline | Primary Strategic Goal |
|--------------------------|---------------------------------------|-----------------------------|---|-----------------|--|
| OpenAI / Stargate | \$500 (plus \$300B in Oracle credits) | 10.0+ | Texas, New Mexico, Ohio, Midwest U.S. | 2025-2028 | AGI Development, U.S. AI Leadership ⁸ |
| Microsoft | \$80 | 2.0+ (Wisconsin alone) | Wisconsin, Iowa, Virginia (U.S.); Global (Europe, Asia, Africa) | Through 2028 | Cloud Dominance, Enterprise AI (Copilot), Vertical Integration ¹² |

| | | | | | |
|----------------------------------|--------------------------|-----------------------------|--|--------------|---|
| Google / Alphabet | \$100+ | 1.5+ (South Carolina alone) | PJM Grid (VA, OH, PA), Iowa, Oklahoma, SC (U.S.); Global | 2025-2027 | Advanced AI Models, Cloud AI Services, Grid Modernization ¹⁴ |
| Amazon Web Services (AWS) | \$75+ (in 2025) | 1.0+ (Indiana alone) | Indiana (U.S.), North America, Europe, Asia Pacific | Ongoing | Cloud AI Dominance, Custom Silicon Efficiency ¹ |
| Meta | Up to \$600 | 5.0+ (Hyperion alone) | Ohio, Louisiana, Kansas (U.S.) | Through 2028 | AGI / "Superintelligence" Development ²⁰ |
| xAI (Colossus) | \$20+ | 1.0+ | Memphis, TN (U.S.) | 2025 | AGI Development (Grok) ² |
| Nvidia | \$100 (in OpenAI) | 10.0 (for OpenAI) | Global (as supplier) | Ongoing | Full-Stack Hardware Dominance ("AI Factory") ²² |
| France (Private Sector) | \$112 | Not specified | France | Ongoing | National/European AI Sovereignty ¹ |
| Total (Announced) | ~\$1.587 Trillion | ~30.5+ GW (Partial) | Primarily U.S. | 2025-2030 | Secure Leadership in AI Era |

1.2 A Planet Re-Wired: The Physical Substrate of Intelligence

The financial figures, while staggering, only tell part of the story. The true significance of this global build-out lies in its physical manifestation: a planetary-scale re-wiring that is creating a new, energy-intensive industrial typology. This is the tangible, material substrate upon which the future of intelligence will be built. According to scaling laws, AI model performance scales with available compute, making this infrastructure build-out a direct race towards AGI.² Current projections estimate the total global AI compute will reach

1027 FLOPS in 2025, a tenfold increase from 2024, enabling models 100 to 1,000 times larger than today's state-of-the-art.²

The most critical metric for understanding this physical transformation is power consumption, measured in gigawatts (GW). A traditional data center might consume 5 to 50 megawatts (MW) of power; the new AI-centric facilities are being designed on a gigawatt scale, requiring up to 2,000 MW (2 GW) each—an increase of two orders of magnitude.²⁷ Aggregating the publicly announced projects reveals a conservative estimate of over 40 GW of new AI-dedicated data center capacity planned or under construction in the United States alone, slated to come online by 2030. This aligns with projections from industry analysts, who estimate that U.S. power demand from AI data centers could surge from 4 GW in 2024 to 123 GW by 2035—a more than thirty-fold increase in just over a decade.²⁷

This immense power demand is not being distributed evenly but is concentrating in a few key geographic hubs, chosen for their access to land, fiber optic networks, and, most importantly, power. Regions like Texas, Ohio, New Mexico, Virginia, and Wisconsin are becoming the epicenters of this new industrial revolution.¹³ This concentration creates unprecedented strain on regional electrical grids, which were not designed to accommodate such large, localized, and constant 24/7 loads. The primary bottleneck to the entire AI revolution is no longer the availability of chips, but the capacity of the grid to deliver power. The current seven-year average wait time for new large-scale projects to secure interconnection to the U.S. grid highlights a fundamental mismatch between the speed of digital innovation and the pace of physical infrastructure development.²⁷

The architecture of these new facilities is also a radical departure from the past. They are not general-purpose data centers but highly specialized "AI Factories," a term explicitly used by Nvidia to describe its integrated hardware and software stack.²² These facilities are physically optimized for the singular task of training and running massive AI models. Their design features high-density racks of GPUs packed so tightly that traditional air cooling is insufficient, necessitating advanced liquid-to-chip or full immersion cooling systems.¹² These racks are interconnected by vast, high-bandwidth networking fabrics, with a single facility containing enough fiber optic cable to circle the Earth multiple times.¹³ This new industrial

typology is engineered for one purpose: the efficient, scaled production of intelligence as a commodity.

Part II: A First-Principles Framework for Analysis: The Holographic Negentropic Imperative

To comprehend the fundamental meaning of the global compute build-out, one must move beyond a purely economic or technological analysis and adopt a framework grounded in the first principles of physics and information theory. The user's provided research offers two such concepts: the Law of Unthinking (LoU) and the Holographic Negentropic Framework (HNF). These frameworks posit that the advance of civilization is not a random historical process but a thermodynamically driven imperative to create order (negentropy) by systematically offloading, or externalizing, complex operations.²

2.1 The Law of Unthinking as a Thermodynamic Driver

The foundational principle for this analysis is the Law of Unthinking (LoU), which elevates a 1911 observation by philosopher and mathematician Alfred North Whitehead into a physical law of progress.² Whitehead wrote, "Civilization advances by extending the number of important operations which we can perform without thinking about them".² He clarified this by comparing the "operations of thought" to "cavalry charges in a battle—strictly limited in number, they require fresh horses, and must only be made at decisive moments".²

This analogy articulates a core biological constraint: conscious cognition is a metabolically expensive and therefore scarce resource. The human brain, despite being only 2% of the body's mass, consumes roughly 20% of its resting energy, equivalent to a continuous power draw of approximately 20 watts.²⁹ This energetic cost frames conscious thought not as an abstract activity, but as a physical, entropy-producing process.

This cognitive constraint is a direct manifestation of the Second Law of Thermodynamics, which states that the entropy (disorder) of a closed system will always increase.⁶ A civilization, however, is not a closed system. It is an open, dissipative structure that maintains and grows its internal complexity by consuming low-entropy energy from its environment and exporting high-entropy waste.² To survive and evolve, such a system must become increasingly efficient

at this process. It must minimize its internal entropy production.³

The Law of Unthinking describes the primary strategy for achieving this thermodynamic efficiency. By offloading a routine or complex operation from the energy-intensive substrate of human cognition onto a more efficient external technology (a tool, a machine, an algorithm), the system conserves its finite "cavalry charges" of conscious thought. This creates a surplus of cognitive and energetic resources that can be reinvested to tackle higher-order problems, leading to the development of even more powerful technologies for offloading. This creates an accelerating, self-reinforcing feedback loop.²

This "Unthinking Advance" can be traced through the major epochs of human history. The Agricultural Era automated energy capture through domesticated animals and gravity-fed irrigation. The Industrial Era automated physical labor with the steam engine and the factory. The Information Era automated rule-based symbolic manipulation with the computer.² The current AI compute boom represents the next, and most profound, stage in this progression: the automation and externalization of cognitive labor itself—of pattern recognition, synthesis, and abstraction. The AI Factory is the literal, industrial-scale embodiment of the Law of Unthinking, designed to mass-produce intelligence as a commodity, thereby making the "operation of thought" an "unthinking" industrial process.²²

2.2 The Holographic Negentropic Framework (HNF) as a Unified Model

The Holographic Negentropic Framework (HNF) provides a more comprehensive, structural model for analyzing this process.³ It synthesizes three foundational pillars from physics and information theory to describe how any complex adaptive system maintains its existence against entropic decay.³

1. **Information Thermodynamics:** This pillar is grounded in Rolf Landauer's principle that "Information is physical".⁶ Every act of information processing, such as erasing a bit of data, has a minimum, unavoidable thermodynamic cost ($kBT\ln 2$).⁶ This principle provides the universal currency—energy and entropy—to measure and compare the efficiency of any system, whether biological or computational.³
2. **The Holographic Principle:** Originating from the study of black hole thermodynamics, this principle posits that the complete description of any three-dimensional volume of space can be thought of as encoded on a two-dimensional boundary surrounding that volume.³⁰ The HNF generalizes this, proposing that all resilient, complex adaptive systems (the 3D "bulk") survive by creating a predictive, error-correcting informational model of their environment on a lower-dimensional sensory or data "boundary." This holographic structure ensures resilience; information is stored in a distributed and

redundant manner, much like a quantum error-correcting code.³

3. **The Law of Unthinking (LoU):** As described above, this is the dynamic engine of the framework. It is the process of performing "negentropic work"—using thermodynamic resources to build and refine the holographic model and to execute efficient, automated actions that maintain the system's internal order.²

When applied to the current global AI build-out, the HNF provides a powerful explanatory model. The entire endeavor can be seen as a civilizational attempt to construct a planetary-scale HNF. The vast, interconnected network of AI data centers acts as the **negentropic regulator**, the engine performing the "unthinking" computational work. The global internet, combined with an exponentially growing array of sensors (satellites, IoT devices, etc.), forms the **holographic boundary**. The data from these sensors is continuously writing information onto this boundary, creating an increasingly high-fidelity informational representation of the physical Earth system. This representation is commonly known as a **Digital Twin of the Earth (DTE)**, a concept that serves as the direct technological manifestation of the HNF's holographic boundary.³ The ultimate purpose of this planetary-scale apparatus, as explicitly stated by many of its architects, is to create a new, higher state of order and intelligence—AGI—capable of modeling, predicting, and ultimately managing the complex dynamics of the world itself.

This framework shares deep parallels with Karl Friston's Free Energy Principle (FEP), which also posits that living systems maintain their integrity by minimizing prediction error (or "surprise") across a statistical boundary known as a Markov blanket.³ The HNF can be seen as a complementary framework that adds a specific structural principle—holography—to the FEP's process-oriented description. It posits that successful, long-lived systems do not just perform inference across a boundary; they evolve a boundary with a resilient, error-correcting, holographic information architecture.³

Part III: The Thermodynamic Ledger: Calculating the Entropic Cost of Intelligence

The construction of this new planetary intelligence substrate, while a monumental act of negentropy (order creation), is subject to the inexorable laws of thermodynamics. The Second Law dictates that this local decrease in entropy must be paid for by a greater increase in the entropy of the surrounding environment.⁶ This "entropic cost" is not an abstract concept but a tangible, measurable externality that manifests as energy consumption, carbon emissions, water usage, and physical waste. This section presents a first-principles calculation of this thermodynamic ledger.

3.1 The Energy Equation: Powering the Global Brain

The most direct entropic cost of the AI build-out is its immense demand for electrical energy. A stark illustration of current inefficiency is the gap between physical limits and reality. According to Landauer's principle, the theoretical minimum energy to perform a bit operation is approximately 2.8×10^{-21} joules at room temperature.⁶ An exaflop-scale AI system performing

10¹⁸ operations per second would thus have a theoretical minimum power draw of just a few milliwatts. In reality, such a system requires on the order of a gigawatt (10⁹ watts)—a gap of roughly 12 orders of magnitude, highlighting the enormous potential for future efficiency gains.²

Calculation Methodology:

The total annual carbon emissions are estimated using the formula:

$$\text{Total CO}_2\text{e} = \sum_{\text{region}} (P_{\text{region}} \times H_{\text{year}} \times CI_{\text{region}})$$

Where:

- P_{region} is the planned AI data center power capacity in a given region (in GW).
- H_{year} is the number of operating hours in a year (8,760).
- CI_{region} is the carbon intensity of the region's electrical grid (in metric tons of CO₂e per GWh).

Input Data and Calculation:

Based on the analysis in Part I, we use a conservative estimate of 40 GW of new AI-dedicated data center capacity coming online in the U.S. by 2030. We will distribute this capacity across the primary build-out regions and apply their respective grid carbon intensities:

- **Texas (ERCOT):** Assumed capacity of 15 GW. The ERCOT grid has a carbon intensity of approximately 389 kgCO₂e/MWh, or 389 metric tons/GWh.
 - Energy Consumption: 15 GW × 8760 h/yr = 131,400 GWh/yr = 131.4 TWh/yr
 - CO₂e Emissions: 131,400 GWh/yr × 389 t/GWh ≈ 51.1 million metric tons/yr
- **Virginia/Ohio/Midwest (PJM & MISO Grids):** Assumed capacity of 15 GW. The PJM and MISO grids, which cover these states, have higher carbon intensities, around 474 kgCO₂e/MWh, or 474 metric tons/GWh.
 - Energy Consumption: 15 GW × 8760 h/yr = 131,400 GWh/yr = 131.4 TWh/yr
 - CO₂e Emissions: 131,400 GWh/yr × 474 t/GWh ≈ 62.3 million metric tons/yr
- **Southwest (WECC Grid - New Mexico, Arizona):** Assumed capacity of 5 GW. The grid in this region has a carbon intensity of approximately 494 kgCO₂e/MWh, or 494 metric tons/GWh.
 - Energy Consumption: 5 GW × 8760 h/yr = 43,800 GWh/yr = 43.8 TWh/yr

- CO₂e Emissions: 43,800 GWh/yr × 494 t/GWh ≈ 21.6 million metric tons/yr
- **Other U.S. Locations:** Assumed capacity of 5 GW at the U.S. average grid intensity.

Projected Output:

Summing these regional estimates, the 40 GW of new AI compute capacity will demand approximately 350 TWh of electricity annually. This is equivalent to nearly 9% of the total U.S. electricity consumption in 2023.³² The associated carbon footprint is projected to be over **135 million metric tons of CO₂e per year**. This massive new load presents a significant challenge to decarbonization goals. While all major tech companies have committed to powering their operations with 100% renewable energy, a fundamental temporal mismatch exists. The exponential growth in compute demand is occurring on a 3-5 year timescale, whereas the transition of the energy grid to renewable sources is a multi-decade project.³⁴ This "sustainability paradox" forces companies to rely on the existing grid, which in key regions remains heavily dependent on fossil fuels. The construction of on-site natural gas power plants at facilities like the Stargate campus in Abilene is a stark admission of this reality: to ensure the required 24/7 reliability, operators must secure firm power, and for now, that often means natural gas.¹¹

3.2 The Water Footprint: Cooling the Engines of Thought

A second, often-overlooked entropic cost is the vast consumption of water for cooling data centers and for the thermoelectric power generation that supplies them.

Calculation Methodology:

The total water footprint is the sum of direct water use for on-site cooling and indirect water use from power generation.

Direct Water Use = $\sum_{\text{region}} (E_{\text{region}} \times WUE_{\text{region}})$

Indirect Water Use = $E_{\text{total}} \times IWC_{\text{grid}}$

Where:

- E_{region} is the annual energy consumption in a region (in kWh).
- WUE_{region} is the Water Usage Effectiveness of data centers in that region (in L/kWh).
- IWC_{grid} is the Indirect Water Consumption factor for the grid's generation mix (in gal/kWh).

Input Data and Calculation:

- **Direct Water Use:** We apply regional WUE values to our energy consumption estimates. WUE can vary dramatically by location and cooling technology, from near zero for

air-cooled systems to over 1.5 L/kWh for evaporative systems in arid regions.³⁷

- Texas (131.4 TWh/yr) with a WUE of 0.24 L/kWh³⁸:
 $131.4 \times 10^9 \text{ kWh} \times 0.24 \text{ L/kWh} \approx 31.5 \text{ billion L/yr} \approx 8.3 \text{ billion gal/yr}$.
- Southwest (43.8 TWh/yr) with a higher WUE of 1.52 L/kWh³⁸:
 $43.8 \times 10^9 \text{ kWh} \times 1.52 \text{ L/kWh} \approx 66.6 \text{ billion L/yr} \approx 17.6 \text{ billion gal/yr}$.
- Using an industry average WUE of 1.9 L/kWh for the remaining capacity⁴, the total direct water consumption for the 40 GW build-out is estimated at **~450 billion gallons annually**. Advanced technologies like closed-loop liquid cooling can reduce this figure by 50-70%, but their deployment is not yet universal.³⁹
- **Indirect Water Use:** Thermoelectric power plants (coal, natural gas, nuclear) withdraw and consume significant amounts of water for steam generation and cooling. The U.S. average is approximately 1.2 gallons per kWh.⁴
 - Total Indirect Water Use: $350 \times 10^9 \text{ kWh/yr} \times 1.2 \text{ gal/kWh} \approx 420 \text{ billion gallons annually}$.

Projected Output:

The combined direct and indirect water footprint of the 40 GW AI build-out is projected to be between 800 billion and 2 trillion gallons of water annually.² This demand places immense pressure on local water resources, particularly in water-stressed regions like the American Southwest, where many of these facilities are being sited.⁴

3.3 The Material Consequence: The E-Waste Cascade

The final entropic cost is the physical matter left behind: a torrent of electronic waste (e-waste). The rapid pace of innovation in AI hardware necessitates aggressive refresh cycles, rendering billions of dollars of equipment obsolete in short order.

Calculation Methodology:

We can estimate the e-waste stream based on the number of servers required, their average weight, and their operational lifespan.

$$\text{E-Waste (tons/yr)} = \frac{\text{Total Servers} \times \text{Avg. Server Weight}}{\text{Avg. Lifespan}}$$

Input Data and Calculation:

A typical 1 GW data center campus requires hundreds of thousands of servers. The Stargate facility in Abilene, for example, will house nearly 500,000 specialized Nvidia chips across its eight buildings.¹¹ Extrapolating to 40 GW suggests a total deployment of 15-20 million servers and specialized AI accelerators. The industry average refresh cycle for data center equipment is 3-5 years to maintain a competitive edge in performance and efficiency.⁴³ Assuming an average server weight of 25 kg and a 4-year lifespan:

$$\frac{17,500,000 \text{ servers} \times 25 \text{ kg/server}}{4 \text{ years}} \approx$$

109,375,000 \text{ kg/yr} \approx 110,000 \text{ metric tons/yr} \$\$

This calculation, based only on servers, is a conservative baseline. A more comprehensive study projects that the rapid expansion of AI could drive e-waste specifically from data centers to as high as 5 million metric tons annually by 2030.⁵ This is a significant contribution to the global e-waste problem, which reached 62 million metric tons in 2022 and is growing five times faster than documented recycling rates.⁴⁴ With only 22.3% of e-waste properly collected and recycled, this new wave of discarded hardware threatens to release toxic materials like lead and mercury into the environment.⁴⁴

The following table summarizes the full entropic cost of the AI compute boom, providing a clear, quantitative ledger of its environmental externalities.

Table 2: The Entropic Cost Ledger—Projected Annual Environmental Impacts of the AI Compute Boom (c. 2030 U.S. 40 GW Build-Out)

| Impact Category | Projected Annual Quantity | Key Assumptions | Contextualization |
|-----------------------------------|-----------------------------|--|---|
| Energy Consumption | ~350 TWh/yr | 40 GW capacity, 24/7 operation | ~9% of 2023 U.S. electricity consumption ³² |
| CO2e Emissions | ~135 Million Metric Tons/yr | Regional grid carbon intensities (TX, PJM, WECC) | Equivalent to the annual emissions of ~30 million gasoline-powered cars |
| Direct Water Consumption | ~450 Billion Gallons/yr | Regional WUEs (0.24-1.9 L/kWh) | Equivalent to the annual water supply for ~4 million U.S. households |
| Indirect Water Consumption | ~420 Billion Gallons/yr | 1.2 gal/kWh for U.S. grid power generation | Equivalent to the annual water supply for ~3.8 million U.S. households |
| E-Waste | 110,000 - | 4-year refresh | Contributes |

| | | | |
|-------------------|--------------------------|-----------------------------|--|
| Generation | 5,000,000 Metric Tons/yr | cycle; external projections | significantly to the 82 million metric tons of global e-waste projected for 2030 ⁴⁴ |
|-------------------|--------------------------|-----------------------------|--|

Part IV: The Negentropic Opportunity: Engineering Environmental Thriving

The thermodynamic ledger presented in Part III quantifies the immense entropic cost of the AI compute boom. From a first-principles perspective, however, this cost is not an argument against the endeavor but rather the necessary investment for an unprecedented negentropic opportunity. The Law of Unthinking and the Holographic Negentropic Framework reveal that this same computational capacity is the essential tool for architecting a new paradigm of planetary stewardship: a transition from reactive "Protection" to proactive "Environmental Thriving."

4.1 From Reactive Protection to Proactive Thriving: A Thermodynamic Critique of Environmental Stewardship

The modern environmental movement and its professional practice were born as a necessary response to the unthinking exploitation of the Industrial Era. This "Protection" paradigm can be characterized as a "conscious brake"—a vast regulatory and administrative apparatus designed to mitigate harm and restrain the entropic outputs of industry.² While essential, this paradigm is fundamentally reactive, problem-focused, and defined by a high-entropy administrative workload of compliance, permitting, and reporting.³

From a thermodynamic perspective, the current process of environmental stewardship is profoundly inefficient. As demonstrated in the user's case study of a TCEQ air permit authorization, the manual workflow consumes vast amounts of its most valuable resource—the cognitive energy of expert engineers—on low-value, automatable "commodity" work like data gathering, calculation, and form-filling.³ This represents a system with high internal entropy (

$S_{\text{manual}} \approx 3.18 \times 10^{-22} \text{ J/K}$) and high informational uncertainty ($H_{\text{manual}} = 2.0 \text{ bits}$), which

requires a large energy input (Emanual=14.4 MJ) to complete.³ It forces the finite "cavalry charges" of expert thought to be squandered on the mundane, rather than being deployed on strategic, high-value challenges.³

The "Agentic Shift"—the application of AI to automate these compliance processes—is the critical first step in a necessary transformation. By applying the Law of Unthinking to its own workflows, the environmental profession can dramatically reduce its internal entropy (Sauto≈1.59×10⁻²² J/K), uncertainty (Hauto≈1.039 bits), and energy cost (Eauto=1.8 MJ).³ This automation is not a threat to the profession; it is a thermodynamic imperative that will generate a massive surplus of cognitive and economic resources.²

This surplus creates the capacity for a new paradigm: "Environmental Thriving".² This emergent model represents a fundamental shift in mindset from reactive to proactive, from fear-based to opportunity-focused. Its goal is not merely to minimize harm but to actively maximize the health, resilience, and biodiversity of planetary systems. It reframes the profession's purpose from managing decline to engineering flourishing. In thermodynamic terms, its objective is to maximize planetary negentropy.⁶

The following table, adapted from the user's provided analysis, outlines the core distinctions between these two paradigms.

Table 3: A Comparative Framework: The "Protection" vs. "Thriving" Paradigms of Environmental Stewardship

| Characteristic | "Protection" Paradigm (Mid-20th Century Model) | "Thriving" Paradigm (Emergent 21st Century+ Model) |
|----------------------------|--|---|
| Core Mindset | Reactive, Problem-focused | Proactive, Solution/Opportunity-focused |
| Primary Goal | Minimize harm, prevent degradation, enforce limits | Maximize systemic health, foster regeneration, cultivate abundance & resilience |
| Dominant Motivation | Fear, anxiety, obligation, guilt | Hope, joy, inspiration, purpose, co-creation |
| Human Role | Steward (as | Co-creator, active |

| | | |
|------------------------------|---|---|
| | controller/corrector of damage) | participant in Earth's negentropic processes |
| Technological Focus | Pollution control, end-of-pipe fixes, monitoring for violations | Information-driven systemic understanding, AI for flourishing |
| Key Metric of Success | Reduction in pollutants, species saved from extinction | Increase in biodiversity, ecosystem vitality, systemic resilience |

4.2 The Environmental General Intelligence (EGI) Hypothesis

The goal of the Thriving paradigm—to manage the entire Earth system for optimal health—is a task of hyper-astronomical complexity, far exceeding the cognitive capacity of any individual human or institution. According to the Law of Unthinking, to make progress on a problem of this scale, its core operations must be automated; they must be made "unthinkable." This requires a new technological substrate, the ultimate expression of which is an **Environmental General Intelligence (EGI)**.²

An EGI is a specialized, planetary-scale AI grounded not in human language and affairs, but in the dynamics of the natural world. Its purpose is not to "think like a person," but to "think like an ecosystem".² The conceptual architecture of such a system, as outlined in the provided research, is a direct application of the Holographic Negentropic Framework³:

- **The Holographic Boundary:** The EGI's sensory input would be a globally integrated network of environmental sensors—satellites, IoT devices, acoustic monitors, eDNA samplers—that continuously feed data into a planetary-scale Digital Twin of the Earth (DTE). This DTE serves as the informational "boundary," encoding the real-time state of the physical biosphere (the "bulk").³
- **The Negentropic Regulator:** The EGI core would be a vast AI system, running on the very compute infrastructure described in Part I, that performs active inference on the DTE. Its function is to build a predictive model of the Earth system, simulate the outcomes of potential interventions, and identify the optimal pathways to guide the planet toward states of higher resilience and health—specifically, keeping it within the safe operating space defined by the Planetary Boundaries framework.³

This vision, while ambitious, is not science fiction. It is the logical integration and scaling of AI

applications that are already being deployed in environmental science today. Researchers are using AI to monitor biodiversity by analyzing satellite imagery and acoustic data, to provide early warnings for wildfires and floods ⁴⁷, to optimize renewable energy grids ⁴⁷, and to create digital twins of natural assets like forests and watersheds to model the impact of conservation efforts. The EGI is the convergence of these disparate efforts into a single, coherent system for planetary stewardship.

4.3 The Thermodynamic Viability of Planetary Intelligence: The Case for More Compute

The creation of a planetary-scale EGI, or "Jed's Angel," can be understood through the lens of the 150-year-old thought experiment of Maxwell's Demon. The demon is an "information engine" that creates a local state of order (negentropy) by acquiring and processing information, at the expense of expending energy and increasing the total entropy of the universe.⁶ The EGI is a real-world instantiation of this concept. Its operation is governed by a strict thermodynamic ledger: the total entropy change of the complete system must be non-negative:

$\Delta S_{\text{Total}} = \Delta S_{\text{EGI}} + \Delta S_{\text{Environment}} \geq 0$.⁶ The system is a net positive only if the value of the created environmental order (

$-\Delta S_{\text{Environment}}$) outweighs the entropic cost of its own operation (ΔS_{EGI}).⁶

This framing reveals a critical, objective truth: the solution to the entropic cost of AI is paradoxically *more* and *smarter* AI. A "stupid" demon that acts inefficiently creates very little order for a high energy cost. An "intelligent" demon, however, can make more precise, targeted interventions, maximizing the negentropic gain for every joule of energy spent. The massive compute build-out, therefore, is the necessary, front-loaded thermodynamic investment to create a more intelligent "demon." This intelligence manifests in several ways:

- **Algorithmic Efficiency:** The process of training ever-larger AI models is an investment of energy to find more ordered and efficient algorithms. This is "burning compute now to save compute later." A more advanced EGI can discover novel methods for climate modeling, materials science, or energy grid optimization that are computationally cheaper to run in the long term, increasing the "intelligence per joule" of our civilization.²
- **Operational Precision:** A more intelligent EGI, equipped with a higher-fidelity Digital Twin of Earth, can perform more precise and effective "negentropic work." Instead of broad, inefficient interventions, it can guide targeted actions—like precision reforestation or the optimized deployment of carbon capture technologies—that achieve the maximum environmental benefit with the minimum energy expenditure and waste.⁶

- **Systemic Resilience:** By creating a more accurate and comprehensive model of the computationally irreducible Earth system, a more powerful EGI enhances our ability to anticipate and mitigate systemic risks like climate tipping points or ecosystem collapse. This proactive management of resilience is a form of negentropy creation that is potentially incalculable in value.³⁰

The physical requirement for more compute is therefore a function of reaching a "thermodynamic breakeven point," where the cumulative negentropic benefit of a highly intelligent planetary regulator begins to decisively outweigh the cumulative entropic cost of its creation and operation.⁶ We are essentially exporting the entropy from our inefficient social and cognitive systems into a more efficient technological substrate, with the goal of achieving a net reduction in the total disorder of the planetary system.²

4.4 A Research Agenda for Planetary-Scale Negentropy

The AI compute boom has placed the environmental profession at a historic crossroads. The immense entropic costs of this build-out represent the greatest new challenge to sustainability, while the computational capacity it provides offers the only tool powerful enough to manage planetary systems at the required scale of complexity. The profession is thus positioned to be either the primary victim of this transition or its primary architect. To navigate this bifurcation point and lead the shift toward the Thriving paradigm, a clear and focused research and development agenda is required:

1. **Embrace the Law of Unthinking Internally:** The profession must accelerate the development and adoption of agentic AI systems to automate the high-entropy work of regulatory compliance and reporting. This is the necessary first step to free up the human capital required for higher-order, creative, and strategic work.²
2. **Harness the New Substrate:** Environmental scientists and engineers must develop the data science and machine learning skills needed to leverage the new planetary compute capacity. This means building and training large-scale environmental models capable of simulating complex ecosystem dynamics and forecasting the impacts of climate change with unprecedented resolution.²
3. **Build the Holographic Boundary:** A concerted, global effort is needed to build the foundational infrastructure for an EGI. This includes prioritizing the development of open-source, interoperable DTE platforms and advocating for the massive expansion of in-situ and remote environmental sensor networks to provide the necessary real-time data.³⁰
4. **Steer the EGI with Wisdom and Equity:** The environmental profession must take a leading role in developing the ethical and governance frameworks for an EGI. This is crucial to ensure its objective functions are aligned with long-term planetary health and

resilience, not with narrow, short-term optimization metrics that could lead to unintended and catastrophic consequences. The risk of misaligned AGI, which some researchers estimate could pose an existential threat, makes this governance challenge paramount.²

Conclusion: The Choice Point—Entropic Collapse or Negentropic Ascent?

The analysis presented in this paper leads to an unambiguous conclusion based on first principles. The unprecedented AI compute build-out of 2025 is the physical manifestation of the Law of Unthinking operating at a planetary scale. It is a thermodynamically driven imperative to externalize and automate the process of intelligence itself. This Great Externalization, while creating the potential for a new, higher state of order (negentropy), carries an immense and immediate environmental price in energy, water, and materials (entropy).

This is not a mere technological trend; it is a fundamental bifurcation point for civilization, a choice point with consequences that will define the coming century. The colossal power of this new computational substrate is a neutral amplifier; its impact will be determined by the goals to which it is applied.

The outcome is not predetermined. If left to proceed under a narrow, unthinking paradigm of pure economic or computational optimization, the entropic costs of the AI revolution could overwhelm planetary systems, accelerating ecological collapse. However, if consciously and deliberately steered, this same computational power provides the necessary, and perhaps only, tool with the requisite complexity to manage the planetary challenges we face. The paradox, from a first-principles analysis, is that the ultimate solution to the problems created by this massive compute build-out is to build an even more intelligent one—an "information engine" so efficient at creating environmental order that its negentropic benefits vastly outweigh its entropic costs.

The environmental profession stands at the fulcrum of this choice. It can remain in its traditional, reactive "Protection" posture and be overwhelmed by a new wave of insurmountable environmental impacts. Or, it can seize this historic opportunity to transform itself, embracing the Law of Unthinking to engineer its own processes and harness this new planetary intelligence. By doing so, it can transition from being a guardian against entropic decay to becoming the proactive architect of a negentropic, thriving planetary system. The Great Externalization presents humanity with a stark choice: between unthinking exploitation leading to systemic collapse, or the deliberate, thoughtful engineering of a sustainable and

intelligent future.

Works cited

1. The Trillion Dollar Horizon: Inside 2025's Already Historic AI ..., accessed September 24, 2025, <https://empirixpartners.com/the-trillion-dollar-horizon/>
2. The Law of Unthinking: An Engine for Environmental Thriving
3. The Thermodynamic Imperative for Automating Environmental Authorizations---Precise Math and Proofs for TCEQ PBR 106.261.docx
4. Data Centers and Water Consumption | Article | EESI - Environmental and Energy Study Institute, accessed September 24, 2025, <https://www.eesi.org/articles/view/data-centers-and-water-consumption>
5. AI-driven data centers risk massive e-waste surge by 2030 - Environmental Health News, accessed September 24, 2025, <https://www.ehn.org/ai-data-center-energy-use>
6. Jed's Angel: A First-Principles Architecture for Planetary Thriving
7. Tech Giants Invest Billions in AI Infrastructure Boom - WebProNews, accessed September 24, 2025, <https://www.webpronews.com/tech-giants-invest-billions-in-ai-infrastructure-boom-2/>
8. OpenAI, Oracle and SoftBank to Build \$500 Billion Stargate AI Data Centers Across U.S., accessed September 24, 2025, <https://www.domain-b.com/technology/technology-general/openai-oracle-and-softbank-to-build-500-billion-stargate-ai-data-centers-across-u-s>
9. Announcing The Stargate Project - OpenAI, accessed September 24, 2025, <https://openai.com/index/announcing-the-stargate-project/>
10. OpenAI, Oracle, and SoftBank expand Stargate with five new AI data ..., accessed September 24, 2025, <https://openai.com/index/five-new-stargate-sites/>
11. OpenAI shows off Stargate AI data center in Texas and plans 5 more elsewhere with Oracle, Softbank, accessed September 24, 2025, <https://apnews.com/article/openai-stargate-oracle-data-center-0b3f4fa6e8d8141b4c143e3e7f41aba1>
12. Microsoft Commits \$80B to AI Data Center Expansion Through 2028, accessed September 24, 2025, <https://www.datacenters.com/news/microsoft-s-80b-investment-in-ai-data-centers-the-digital-backbone-for-a-multimodal-world>
13. Made in Wisconsin: The world's most powerful AI datacenter - Microsoft On the Issues, accessed September 24, 2025, <https://blogs.microsoft.com/on-the-issues/2025/09/18/made-in-wisconsin-the-worlds-most-powerful-ai-datacenter/>
14. Google commits to \$25 billion investment in AI infrastructure and ..., accessed September 24, 2025, <https://www.mitrade.com/insights/news/live-news/article-3-960558-20250715>
15. Google Commits \$25 Billion to AI and Data Center Expansion Across the U.S.'s Largest Electric Grid, accessed September 24, 2025,

- <https://odsc.medium.com/google-commits-25-billion-to-ai-and-data-center-expansion-across-the-u-s-s-largest-electric-grid-52540397623d>
16. Google Cloud to pour more than \$25B into AI infrastructure across PJM - Utility Dive, accessed September 24, 2025, <https://www.utilitydive.com/news/google-cloud-blackstone-aws-us-ai-data-center-buildouts/753202/>
 17. Data Centers - AWS Sustainability, accessed September 24, 2025, <https://aws.amazon.com/sustainability/data-centers/>
 18. Artificial Intelligence (AI) on AWS - AI Technology, accessed September 24, 2025, <https://aws.amazon.com/ai/>
 19. AI Infrastructure on AWS – Artificial Intelligence Innovation Capabilities, accessed September 24, 2025, <https://aws.amazon.com/ai/infrastructure/>
 20. Mark Zuckerberg Warns of Risk in 'Mispending Billions' Chasing AI, but Calls Underinvestment a Greater Danger - MLQ.ai | Stocks, accessed September 24, 2025, <https://mlq.ai/news/mark-zuckerberg-warns-of-risk-in-mispending-billions-chasing-ai-but-calls-underinvestment-a-greater-danger/>
 21. Meta Unveils Billion-Dollar AI Data Centre Push - Digit.fyi, accessed September 24, 2025, <https://www.digit.fyi/meta-ai-investment/>
 22. AI Factories Are Redefining Data Centers, Enabling Next Era of AI | NVIDIA Blog, accessed September 24, 2025, <https://blogs.nvidia.com/blog/ai-factory/>
 23. Nvidia to invest \$100 billion in OpenAI to help expand ChatGPT ..., accessed September 24, 2025, <https://apnews.com/article/openai-nvidia-investment-partnership-chatgpt-610d894d93f9be23c46762950997a67f>
 24. More questions than answers in Nvidia's \$100 billion OpenAI deal, accessed September 24, 2025, <https://indianexpress.com/article/technology/tech-news-technology/more-questions-than-answers-in-nvidias-100-billion-openai-deal-10266666/>
 25. Understanding Microsoft Datacenters, accessed September 24, 2025, <https://news.microsoft.com/datacenters/>
 26. Investing in America 2025 - Google Blog, accessed September 24, 2025, <https://blog.google/inside-google/company-announcements/investing-in-america-2025/>
 27. AI infrastructure gaps | Deloitte Insights, accessed September 24, 2025, <https://www.deloitte.com/us/en/insights/industry/power-and-utilities/data-center-infrastructure-artificial-intelligence.html>
 28. Top 10 Digital Infrastructure Projects to Watch in 2025 - 174 Power Global, accessed September 24, 2025, <https://174powerglobal.com/blog/top-digital-infrastructure-projects-to-watch/>
 29. Inverting the Stack: Environmental Intelligence
 30. The Simplicity Imperative: A Unified Framework for Information, Computation, and Planetary Stewardship
 31. The Thermodynamic Ledger of the Cosmos: From Black Hole Information to Planetary Thriving

32. Data Centers and Their Energy Consumption: Frequently Asked Questions - Congress.gov, accessed September 24, 2025, https://www.congress.gov/crs_external_products/R/PDF/R48646/R48646.1.pdf
33. Data Centers and Their Energy Consumption: Frequently Asked Questions - EveryCRSReport.com, accessed September 24, 2025, <https://www.everycrsreport.com/reports/R48646.html>
34. Is AI's energy use a big problem for climate change?, accessed September 24, 2025, <https://climate.mit.edu/ask-mit/ais-energy-use-big-problem-climate-change>
35. AI: Five charts that put data-centre energy use – and emissions – into context - Carbon Brief, accessed September 24, 2025, <https://www.carbonbrief.org/ai-five-charts-that-put-data-centre-energy-use-and-emissions-into-context/>
36. As generative AI asks for more power, data centers seek more reliable, cleaner energy solutions - Deloitte, accessed September 24, 2025, <https://www.deloitte.com/us/en/insights/industry/technology/technology-media-and-telecom-predictions/2025/genai-power-consumption-creates-need-for-more-sustainable-data-centers.html>
37. What Is Water Usage Effectiveness (WUE) in Data Centers? - The Equinix Blog, accessed September 24, 2025, <https://blog.equinix.com/blog/2024/11/13/what-is-water-usage-effectiveness-wue-in-data-centers/>
38. Measuring energy and water efficiency for Microsoft datacenters, accessed September 24, 2025, <https://datacenters.microsoft.com/sustainability/efficiency/>
39. Optimizing water usage effectiveness for data centers - Vertiv, accessed September 24, 2025, <https://www.vertiv.com/en-cn/about/news-and-insights/articles/educational-articles/optimizing-water-usage-effectiveness-for-data-centers/>
40. Circular water solutions key to sustainable data centres - The World Economic Forum, accessed September 24, 2025, <https://www.weforum.org/stories/2024/11/circular-water-solutions-sustainable-data-centres/>
41. The world's AI generators: rethinking water usage in data centers to build a more sustainable future - Lenovo StoryHub, accessed September 24, 2025, <https://news.lenovo.com/data-centers-worlds-ai-generators-water-usage/>
42. The Negentropic Channel--A First-Principles.pdf
43. Data Centers The Environment - Supermicro, accessed September 24, 2025, https://www.supermicro.com/wekeepitgreen/Data_Centers_and_the_Environment_Dec2018_Final.pdf
44. The global E-waste Monitor 2024 – Electronic Waste Rising Five Times Faster than Documented E-waste Recycling: UN, accessed September 24, 2025, <https://ewastemonitor.info/the-global-e-waste-monitor-2024/>
45. Global e-Waste Monitor 2024: Electronic Waste Rising Five Times Faster than Documented E-waste Recycling | UNITAR, accessed September 24, 2025, <https://unitar.org/about/news-stories/press/global-e-waste-monitor-2024-electro>

- [nic-waste-rising-five-times-faster-documented-e-waste-recycling](#)
46. Data Center Recycling: Limits of E-Waste Recycling Solutions - Human-I-T, accessed September 24, 2025, <https://www.human-i-t.org/data-center-recycling/>
 47. 10 Real Examples of Sustainable AI Transforming Planet, accessed September 24, 2025, <https://www.sentisight.ai/10-real-life-examples-sustainable-ai-action/>
 48. Top 10 Sustainability AI Applications & Examples - Research AIMultiple, accessed September 24, 2025, <https://research.aimultiple.com/sustainability-ai/>
 49. AI Technology is Revolutionizing Climate Change Mitigation - Appen, accessed September 24, 2025, <https://www.appen.com/blog/how-ai-technology-is-revolutionizing-climate-change-mitigation>
 50. How AI can help mitigate climate change and drive business efficiency | Carbon Direct, accessed September 24, 2025, <https://www.carbon-direct.com/insights/how-ai-can-help-mitigate-climate-change-and-drive-business-efficiency>
 51. AI and environmental challenges | UPenn EII, accessed September 24, 2025, <https://environment.upenn.edu/news-events/news/ai-and-environmental-challenges>