Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

The Hydro-Digital Paradox: Water Scarcity in the Age of Artificial Intelligence

Aishwarya Natarajan
Amazon Web Services, USA
Corresponding Author: Aishwarya Natarajan, **Email**: aishwarya.mamandur@gmail.com

| ABSTRACT

The development of artificial intelligence is rapidly accelerating, and the demand for computational infrastructure is similarly growing. While these rapidly expanding data centers, known as hyperscale data centers, are of crucial importance for both instruction and operation of advanced AI models, the environmental consequences of this water consumption remain largely invisible to end users. While data centers remain a small fraction of total water consumption, increasing AI capabilities have led to massive increases in water consumption, in many cases, amounting to water consumption similar to that of thousands of households. Hyperscale facilities are always located in urban areas and often drought areas, which creates a degree of competition for scarce resources while raising questions of environmental concerns and community well-being. From a structural standpoint, the cooling of AI infrastructure has significant water requirements and is routinely accomplished via evaporative cooling, which represents one of the most water-intensive activities in the digital economy. As AI becomes more and more integrated into the global system--economically and socially--the issues raised by the need for physical infrastructure that sustains this activity are an urgent challenge for policymakers, technology, and environmental advocates.

| KEYWORDS

Artificial intelligence, water consumption, data centers, environmental sustainability, resource scarcity

| ARTICLE INFORMATION

ACCEPTED: 12 June 2025 **PUBLISHED:** 10 July 2025 **DOI:** 10.32996/jcsts.2025.7.7.55

1. Introduction: The Hidden Environmental Cost of Artificial Intelligence

The Contradiction Between AI Progress and Ecological Preservation

The advent of artificial intelligence now poses a unique challenge for environmental scientists and policy-makers [1]. Although AI is presenting potential solutions to climate modeling, resource allocation, and others, the physical infrastructure that supports these systems uses resources (and generates waste) at historically unprecedented rates. In the case of large language models, the training alone requires server farms (which can run continuously for several weeks) and will, in turn, produce substantial heat loads that conventional refrigeration systems cannot adequately manage. To expand their hi-tech endeavors, data centers have begun adopting water-intensive cooling strategies. This practice comes at a time when many parts of the world are experiencing historically unprecedented droughts and water restrictions. Ironically, after decades of technologies that enable more efficient resource use, it now appears that relatively traditional technologies are re-consuming substantially more resources [2]. This contradiction requires us to think differently about how society measures technological advancement, particularly when the technologies intended to alleviate socio-environmental challenges could only exacerbate them further.

Overview of AI's Growing Environmental Footprint

Behind every AI interaction lies a physical reality of servers, cooling towers, and water pumps working around the clock [1]. Modern data centers employ evaporative cooling as their primary temperature control method, a process that transforms freshwater into vapor and releases it into the atmosphere. Unlike closed-loop systems that recirculate coolant, evaporative cooling permanently removes water from local watersheds. Water usage has surged dramatically alongside the increasing sophistication and size of artificial intelligence systems. Training runs for advanced language models can span months, during which cooling systems operate continuously at maximum capacity [2]. Tech companies have clustered their facilities near urban centers to reduce latency, but these locations often coincide with regions already struggling to meet residential and agricultural water needs. The result is a new form of resource competition that pits digital infrastructure against traditional water users.

Thesis Statement

The emergence of AI data centers in water-scarce areas is driving a wedge between technological aspiration and reality on the ground [1][2]. Communities from Arizona to Northern Virginia are witnessing the creation of large facilities owned by tech companies, consuming water resources equivalent to that of a small city, as surrounding local aquifers decline and also have tighter drought restrictions. These facilities mark more than simply a development in infrastructure; they showcase a fundamental disconnect between where we need computational power and where there is a water resource. Therefore, the current trajectory suggests the risk that, absent intervention, AI development will exacerbate water stress in already vulnerable communities, ultimately turning a tool of advancement into an instrument of environmental injustice. Responding to this challenge involves not only technical solutions but also a consideration of where and at what scale to build AI infrastructure.

2. Quantifying the Thirst: Water Consumption Metrics of Major Tech Companies

Case Study: Data Center Operations in The Dalles, Oregon

A major technology corporation's facility in The Dalles represents a critical example of how leading tech companies have transformed small communities into hubs for computational infrastructure [3]. The site selection in Oregon took advantage of nearby hydroelectric dams and the region's moderate temperatures, yet this strategic positioning has created unexpected tensions with existing water users. What began as promises of low environmental impact has evolved into significant water withdrawals from the Columbia River basin, often occurring when farmers in the area face irrigation cutbacks. Local officials have struggled to balance the economic benefits of hosting such facilities against mounting concerns from residents who question whether digital infrastructure should take precedence over traditional water uses [4]. The situation in The Dalles has become a reference point for other communities considering similar developments, highlighting the need for tech companies to transparently report resource consumption.

Escalating Water Usage Among Industry Leaders

Leading technology corporations' global data center networks reveal troubling patterns of increasing water demand that correlate directly with expansion into artificial intelligence services [3]. Annual environmental reports from these companies celebrate their renewable electricity purchases while quietly documenting sharp rises in water use at facility after facility. Each new AI feature launch and every language model upgrade translates into higher cooling demands and more water consumed. The speed at which water needs multiply once companies commit to AI development has caught many observers off guard [4]. The experience of major tech companies serves as a cautionary tale for other corporations planning similar technological transitions without first securing sustainable cooling alternatives.

Metric	Description	Scale of Impact
Annual Water Usage	Total consumption by major facilities	Millions to billions of gallons
Peak Daily Consumption	Maximum usage during intensive operations	Hundreds of thousands of gallons
Year-over-Year Growth	Rate of increase in water consumption	Double-digit percentage increases
Comparison to Residential Use	Equivalence to household consumption	Thousands of homes

Table 1: Water Consumption Patterns of Major Data Center Operations [3, 4]

Peak Consumption Projections

Major social media companies' data center operations exemplify the challenge of predicting and managing water consumption in facilities designed for variable computational loads [3]. Their infrastructure must accommodate both baseline platform operations and intensive AI model training, creating peaks in resource demand that strain local water systems. During periods of maximum computational activity, these facilities require cooling capacity that can overwhelm municipal water supplies, particularly in regions already experiencing seasonal variations in availability. The unpredictable nature of these consumption spikes complicates planning efforts by both corporations and host communities [4]. This situation illustrates how the intermittent but intense water demands of AI development can create more disruption than steady industrial consumption patterns.

Comparative Analysis: AI Training Water Consumption vs. Traditional Manufacturing

The water footprint of artificial intelligence training presents a stark contrast to traditional manufacturing processes in both scale and efficiency [3]. While the automotive and textile industries have developed water recycling systems over decades of environmental regulation, data centers often rely on single-use evaporative cooling that provides no opportunity for water recovery. The comparison becomes particularly troubling when considering the output: manufacturing facilities produce tangible goods with clear utility, while AI training runs may generate models that never reach deployment or quickly become obsolete. Research indicates that training a single large language model can consume water quantities comparable to manufacturing hundreds of vehicles, yet without the established regulatory frameworks that govern traditional industry [4]. This disparity highlights the urgent need for water consumption standards specifically tailored to computational infrastructure.

3. The Technical Imperative: Understanding AI's Cooling Requirements

Computational Demands of Large Language Model Training

Training contemporary language models requires computational operations that push hardware to its absolute limits for extended periods [5]. Each training cycle involves processing vast datasets through billions of parameters, with graphics processing units and specialized chips operating at maximum capacity throughout. These marathon sessions generate thermal loads that would destroy components without aggressive cooling intervention. The shift from traditional computing workloads to AI training has fundamentally altered data center design requirements, as facilities must now handle sustained peak performance rather than variable loads. Engineers face the challenge of maintaining stable temperatures across thousands of processors working in parallel, where even minor thermal variations can cause training failures or hardware damage [6]. The continuous nature of model training leaves no downtime for cooling systems to recover, creating unprecedented stress on infrastructure originally designed for intermittent computational spikes.

Heat Generation in Hyperscale Data Centers

Modern data facilities housing AI infrastructure produce heat densities that exceed traditional server rooms by orders of magnitude [5]. Rack-mounted systems configured for machine learning pack processors so densely that air cooling alone cannot prevent thermal runaway. The concentration of computational power in these facilities creates hot spots where temperatures would rapidly reach component failure thresholds without intervention. Unlike conventional data centers, where heat generation varies with user demand, AI training facilities maintain a constant high thermal output for weeks or months at a time. This sustained heat production challenges fundamental assumptions about data center cooling design, forcing engineers to implement solutions borrowed from industrial processes [6]. The thermal density problem compounds as facilities expand, creating scenarios where heat removal becomes the primary constraint on computational capacity.

Evaporative Cooling Systems and Their Water-Intensive Nature

Evaporative cooling has emerged as the dominant solution for managing extreme heat loads in AI facilities, despite its significant water consumption [5]. This approach works by passing warm air over water-saturated media, using evaporation to absorb heat before exhausting the now-humid air outside. The process offers superior cooling efficiency compared to traditional air conditioning but permanently removes water from local supplies through evaporation. Data centers have adopted industrial-scale cooling towers similar to those used in power plants, transforming facilities into major water consumers in their communities. The technology's effectiveness at managing high heat loads makes it nearly irreplaceable for current AI workloads, even as water scarcity concerns mount [6]. Alternative cooling methods exist but require either proximity to specific geographic features or investments that many operators consider prohibitive.

Cooling Method	Water Consumption Level	Efficiency Rating	Environmental Impact
Evaporative Cooling	Very High	High efficiency	Permanent water loss
Air Cooling	None	Low efficiency	High energy use
Liquid Cooling	Moderate	Moderate efficiency	Potential for recycling
Hybrid Systems	Variable	Variable	Mixed impact

Table 2: Cooling Technologies and Resource Requirements [5, 6]

The Relationship Between AI Model Complexity and Resource Consumption

A direct correlation exists between model sophistication and cooling requirements, with each generation of AI systems demanding exponentially more thermal management [5]. As researchers push toward more capable models with increased parameters and deeper architectures, the computational intensity rises correspondingly. This relationship creates a feedback loop where advances in AI capability directly translate to greater environmental resource demands. The pursuit of marginal improvements in model performance can require doubling or tripling computational resources, with proportional increases in cooling needs. Industry trends suggest this trajectory will continue as competition drives the development of ever-larger models [6]. The technical community faces a fundamental question about whether the benefits of slightly improved AI performance justify the escalating environmental costs of achieving those gains.

4. Geographic Convergence: Water Stress and Data Center Placement

Strategic Location of Data Centers Near Urban and Industrial Hubs

Data center developers consistently select sites adjacent to metropolitan areas and manufacturing zones, driven by requirements for minimal network latency and access to skilled technicians [7]. These locations offer existing fiber optic infrastructure and proximity to business customers who demand rapid response times for cloud services. However, urban and industrial regions typically face higher water demand from established users, creating immediate competition for limited resources. The convergence of digital infrastructure with traditional urban water systems has produced unforeseen strain on municipal supplies designed for residential and commercial use. Local governments find themselves mediating between technology companies promising economic development and existing water users concerned about resource security [8]. This geographic clustering pattern repeats globally, as data center operators prioritize connectivity over water availability in their site selection process.

Regional Analysis: Arizona, Northern Virginia, and India

Desert regions like Arizona have become unexpected hotspots for data center development, despite obvious water scarcity challenges [7]. The state's favorable tax policies and reliable power grid attract facilities that then compete with agriculture and urban users for Colorado River allocations. Northern Virginia's data center concentration ranks among the world's highest, even as groundwater levels sink year after year and dry periods stretch longer each summer. India faces a different puzzle - the monsoon dumps most yearly rainfall in just a few months, leaving cities and tech facilities scrambling for stored water through the long dry season. Each region demonstrates how economic incentives can override environmental constraints in facility placement decisions [8]. The pattern reveals a systematic misalignment between where computational infrastructure is located and where sustainable water resources exist.

Region	Water Stress Level	Data Center Concentration	Primary Water Source	Competing Uses
Arizona	Severe	High	Colorado River	Agriculture, Urban
Northern Virginia	Moderate to High	Very High	Aquifers	Residential, Commercial
India	Seasonal/Variable	Growing	Monsoon-dependent	Urban, Agriculture
Oregon	Low to Moderate	High	River systems	Agriculture, Industry

Table 3: Geographic Distribution of Water Stress and Data Center Density [7, 8]

Conflict Between Existing Drought Conditions and New Infrastructure Demands

Many data center developments proceed in regions already experiencing multi-year droughts and declining water tables [7]. Site planners often reference precipitation records from the 1990s or earlier, building for weather patterns that disappeared years ago. Meanwhile, concrete trucks roll past farm fields where irrigation has stopped and residential areas under strict conservation orders. When facilities finally power up, they often discover the water picture has changed dramatically from what their feasibility reports promised. Resident groups now regularly pack town halls to oppose new data centers, having learned that these facilities keep their cooling towers running no matter how severe the drought [8]. The gap between corporate timelines and environmental reality demands new approaches to measuring water availability that look forward, not backward.

Local Water Resource Competition and Allocation Challenges

Water allocation frameworks developed for traditional users struggle to accommodate the unique consumption patterns of data centers [7]. Agriculture has flexibility - skip watering, delay planting, and let some fields fallow. Residents adapt too, through low-flow fixtures and xeriscaped yards. But servers crash if temperatures rise, making data center water needs absolutely inflexible. Water boards find themselves in impossible situations, weighing bytes against bushels, cooling towers against kitchen taps. Existing water laws rarely mention digital infrastructure, which was written when computers filled entire rooms rather than entire districts [8]. Traditional water rights systems face unprecedented challenges in determining whether computational services deserve equal priority with food production and human consumption.

5. Community Stakes: When Digital Progress Meets Local Water Rights

Trading Tax Revenue for Water Security

Communities that welcome data centers often do so with promises of job creation and tax revenue ringing in their ears [9]. Local politicians tout the arrival of tech facilities as validation of their region's connectivity and business climate. Yet the jobs created tend toward security and maintenance roles rather than the high-tech positions communities envision. Property tax revenues, while substantial, rarely offset the infrastructure upgrades required to support massive water and power demands. The economic calculus becomes even murkier when factoring in opportunity costs - water allocated to cooling servers cannot irrigate crops or fill swimming pools. Some communities have discovered that hosting data centers means accepting a permanent industrial presence that operates around the clock, transforming the character of formerly rural areas [10]. The promised economic transformation often amounts to a handful of jobs and tax payments that pale beside the environmental burden imposed.

Stakeholder Group	Potential Benefits	Environmental Costs	Net Impact Assessment
Local Communities	Tax revenue, Limited jobs	Water depletion, Infrastructure strain	Generally negative
Tech Companies	Operational efficiency	Minimal direct cost	Highly positive
Agricultural Users	None	Water competition	Negative
Future Generations	Technology access	Depleted resources	Severely negative

Table 4: Stakeholder Impact Assessment [9, 10]

Public Concern Over Resource Depletion

Local pushback keeps growing stronger each time another facility gets proposed, especially after summers when wells start failing [9]. What used to be quick zoning approvals turned into marathon hearings packed with farmers, homeowners, and small business operators asking hard questions. Neighborhood Facebook groups circulate photos - dust-dry canals on one side of the road, misting towers consuming thousands of gallons per minute on the other. Environmental groups that spent decades fighting chemical plants now train their sights on server farms, recognizing the threat runs just as deep. The promise of technological progress rings hollow for families facing water cutoffs while watching vapor plumes rise from nearby server facilities [10]. Public sentiment has shifted from welcoming tech investment to questioning whether any amount of tax revenue justifies compromising water security for future generations.

Environmental Justice Considerations in Water-Scarce Regions

Data center placement often follows patterns that burden already vulnerable communities with additional environmental stress [9]. Poor neighborhoods and farm towns end up hosting the most facilities - land sells cheap there, and organized pushback stays minimal. Most residents in these places don't have lawyers on speed dial or funds to hire environmental consultants who might spot problems in permit applications. The benefits of AI advancement flow primarily to urban professionals and tech companies, while rural and minority communities bear the cost through depleted aquifers and competition for water access. Indigenous communities with traditional water rights find themselves competing against corporate users with deep pockets and political connections [10]. This geographic inequality in who benefits from AI versus who pays its environmental price raises fundamental questions about fairness in the digital economy.

Policy Challenges in Regulating Data Center Water Usage

Regulatory frameworks struggle to keep pace with the rapid expansion of data center infrastructure and its unique water demands [9]. Traditional zoning laws and water permits never anticipated facilities that might consume as much water as entire towns while employing only dozens of workers. Policymakers face lobbying from tech companies threatening to locate facilities elsewhere if restrictions prove too stringent. The global nature of digital services complicates local regulation - should a rural county restrict water use for facilities serving customers worldwide? Tech lobbyists warn that recycling mandates or hard consumption caps would force them to build in Mexico or Canada instead, taking jobs and tax dollars with them [10]. Officials fear that getting tough locally just means someone else's aquifer gets drained, solving nothing while losing economic opportunities.

Conclusion

The explosive growth of artificial intelligence infrastructure has placed a temporal collision between miraculous technological advancement and environmental sustainability, with water resources bearing the largest burden of a growing crisis. AI-relevant data centers have become massive water consumers in regions that are already receding into water scarcity, creating uncertainty that threatens the very resilience of impacted communities. Whether considering the geographic disparity between computational infrastructures and a viable water source and learning the inflexible cooling requirements of water-consuming structures, these trends establish a technocentric reality where developing technological systems actively deteriorate or halt local water security. Communities that are actively hosting these structures are being faced with uncomfortable trade-offs between economic development and the preservation of accessible resources, while regulatory bodies are faced with unique challenges due to the data centers, over which they have little control. The environmental justice issues are even more pronounced as many of the resources in water-scarce areas, particularly water, are burdened by the vulnerable and historically marginalized members of society who ultimately receive the least benefits of AI advancement. If the ideals of engineering, site selection, and regulation do not change, the present trajectories of the AI industry will worsen the current global crisis of water availability. The way forward must include a reimagining of how society builds and operates housing for artificial intelligence, and that system of calculation does not simply rely on computational metrics when managing water resources that are necessary for human beings or vital life processes.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.
Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

[1] Arti Garg, et al., "The Hidden Cost of AI: Unpacking its Energy and Water Footprint." IEEE Technology Climate Center (ITCC), 25 February 2025. https://itcc.ieee.org/blog/the-hidden-cost-of-ai-unpacking-its-energy-and-water-footprint/
[2] IEEE Computer Society Team, "AI and Resource Consumption: Examining the Environmental Impact." 18 February 2025. https://www.computer.org/publications/tech-news/research/ai-and-the-environment
[3] Almando Morain, et al., "Artificial Intelligence for Water Consumption Assessment: State of the Art Review." Water Resources Management, 25 April 2024. https://link.springer.com/article/10.1007/s11269-024-03823-x
[4] HyScaler. "Water Consumption of AI: How Tech Giants are Draining the Planet." 26 February 2024. https://hyscaler.com/insights/water-consumption-of-ai-tech-giants/
[5] Samuel K. Moore, "Tiny 'Fans-on-Chips' Could Cool Big Data Centers > xMEMS Makes Tiny Coolers for Big Data Centers." IEEE Spectrum, 2 May 2025. https://spectrum.ieee.org/data-center-cooling-xmems
[6] Ali Heydari. "Keynote 3: How Generative AI and Accelerated Computing are Creating the Next Generation Liquid Cooled Data Centers with Focus on Challenges, Opportunities, and the Road Ahead." IEEE ITherm Conference, 2024. https://www.ieee-itherm.net/keynote-3-2024/
[7] Mohammad A. Islam, et al., "Water-Constrained Geographic Load Balancing in Data Centers." IEEE Transactions on Cloud Computing, Volume 5, Issue 2, July 8, 2015. https://nweb.eng.fiu.edu/gaquan/Papers/islam2015TCC.pdf
[8] Alex Setmajer, "How Data Centers Use Water, and How We're Working to Use Water Responsibly." Interconnections - The Equinix Blog, September 19, 2024. https://blog.equinix.com/blog/2024/09/19/how-data-centers-use-water-and-how-were-working-to-use-water-responsibly/
[9] Fanchao Yu, "Sustainable Development and Economics: Balancing Growth with Environmental Concerns." Business and Economics Journal, Volume 14, Issue 4, 2023. https://www.hilarispublisher.com/open-access/sustainable-development-and-economics-balancing-growth-with-environmental-concerns.pdf
[10] Shuo Vivien, "Sustainable Development: Balancing Economic Prosperity and Environmental Concerns." Academy of Business Research Journal, 2023. https://www.abacademies.org/articles/sustainable-development-balancing-economic-prosperity-and-environmental-concerns.pdf