

MYTHBUSTER

10 myths about AI data centers and water.

And what the numbers
actually say.




Table of contents

[01]	Myth 1: “Every ChatGPT conversation drinks a bottle of water”
[02]	Myth 2: “All data centers cool with water”
[03]	Myth 3: “Water used by a data center is gone forever”
[04]	Myth 4: “AI’s water footprint is mostly at the data center”
[05]	Myth 5: “Any water use is bad water use”
[06]	Myth 6: “Data centers are uniquely enormous water users”
[07]	Myth 7: “More compute always means more water”
[08]	Myth 8: “Data centers take drinking water from communities”
[09]	Myth 9: “Water use is just the cost of progress, and nothing can be done”
[10]	Myth 10: “The numbers can be taken at face value”
[11]	The real questions to ask
[12]	Sources and boundaries

Introduction

Incorrect information and lazy assumptions, whether from the media or the data center developers themselves, create real and lasting harm for the AI industry and the communities that host it.


“Every ChatGPT conversation drinks a bottle of water.” It is the single most famous statistic in the debate over AI’s environmental cost. The image is vivid, portable, and almost perfectly designed to go viral. It has shaped public and political conversation about AI infrastructure more than any other single statistic.

The concern behind it is legitimate. Data centers are being built at extraordinary speed, some of them in water-stressed regions, and much of the installed base was designed in an era when water efficiency was not a primary design consideration; communities have in several documented cases had genuine cause for grievance. But the public debate is now running on figures that range from outdated to wrong, and wrong on many levels. Some of the loudest claims overstate the problem by orders of magnitude. Others, including several the industry itself repeats, understate it.

Getting the numbers right matters most to the communities that host these facilities, because bad numbers produce bad siting decisions, misdirected scrutiny, and misplaced trust. Here are ten myths about AI and water, with the sources for each, concluded with honest caveats on where the corrections apply and where they do not.

What this paper covers

- 10 myths about AI and water
- Submer’s considered response to each myth
- Data sources for the myths and responses
- Caveats to consider



[MYTH 1]

“Every ChatGPT
conversation
drinks a bottle
of water”

10 MYTHS ABOUT
AI DATA CENTERS
AND WATER

[MYTH 1]

“Every ChatGPT conversation drinks a bottle of water”

One headline.
Three different
boundaries.

The claim traces to a single academic paper: [Making AI Less “Thirsty”](#) by Pengfei Li, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren of UC Riverside, first posted in 2023 and later published in Communications of the ACM. The researchers estimated that GPT-3 consumed roughly 500ml of water per 10 to 50 medium-length responses, not per query, and their figure was a calculated estimate built from public data, not a meter reading from inside a data center. [A 2024 Washington Post analysis](#), produced with Ren, then restated the number as roughly a bottle of water per 100-word GPT-4 email, and that version is the one that went viral.

Two things happened next that the viral version never caught up with. First, Ren himself has since [indicated](#) the GPT-4 estimate was likely far too high and that the real figure is closer to 15ml per prompt in total, of which around 5ml is on-site cooling water. Second, in August 2025 [Google](#)

[published the first large-scale measurement](#) from a production environment, finding the median Gemini text prompt consumed 0.26ml of on-site water, or about five drops.

In summary, the “bottle per query” framing is not supported by anyone, including its original author. The five-drops figure is also not the whole story. It counts only on-site cooling water and excludes the water consumed generating the electricity that powers the data center, which, as Myth 4 shows, is usually the larger number. Credible per-prompt estimates still span more than a thousandfold depending on model, location, season, and, above all, what you choose to count. The honest answer to “how much water does a query use?” is a range, and the range is wide because disclosure is inconsistent, a problem we return to in Myth 10. Any single confident number, large or small, deserves scrutiny before it deserves a headline.

What the numbers say	ORIGINAL ESTIMATE	REVISED ESTIMATE	ON-SITE WATER ONLY
	500 ml	15 ml	0.26 ml
	per 10–50 GPT-3 responses	per prompt, total water	per median Gemini prompt
THE CAVEAT			
The three figures are not directly comparable. They refer to different models, workloads, methods, and accounting boundaries			

Grouped sources for this section can be found on [page 28](#), section [01]

[MYTH 2]

“All data centers cool with water”

10 MYTHS ABOUT
AI DATA CENTERS
AND WATER

[MYTH 2]

“All data centers cool with water”

Cooling architecture changes everything.

Data center cooling systems comprise primary and secondary elements. The primary element takes care of the heat in the data hall, either via air or liquid. Secondary cooling is the element of the system that rejects heat outside the data center, and can come in several different technological variants:



Evaporative cooling

[01]

Evaporative cooling sprays or trickles water through towers and loses much of it to the atmosphere. This is the “thirsty” method, and the source of most legitimate concern, although modern-day systems are increasingly efficient and can help displace indirect water consumption in certain grid regions (see Myth 5).



Chillers

[02]

Chillers are among the most common technology deployed in new data center developments and standard models consume no water at all, using a process of compression and expansion of refrigerant to absorb and reject heat via heat exchangers. Some chillers are fitted with adiabatic systems that use water mist or wetted media to pre-cool the ambient air during hot summer peaks before it reaches the chiller’s condenser coils, but deployments of this type are relatively rare, do not operate all year round, and are quite water-efficient when they do.



Dry coolers

[03]

Dry coolers reject heat directly to ambient air across a heat exchanger, using no water and no mechanical refrigeration. These are the most efficient option where conditions allow, although they are limited by data hall operating temperatures and ambient conditions. Hybrid versions with chiller functionality are commonly used in modern designs, allowing for use in a wider temperature range.

[MYTH 2]

“All data centers cool with water”

Closed loops recirculate the **same water indefinitely.**

Any system involving chillers and/or dry coolers receives heat from the data halls via water, but in a closed loop system. These sealed water systems are filled once during commissioning and recirculate the same water indefinitely.

In the data center itself we find the primary cooling system, and technology choices in this area are evolving due to increased heat loads associated with the processing of compute-intensive IT workloads like AI. This means that liquid (generally a dielectric fluid or mix of water and propylene glycol) is becoming the only viable cooling medium, replacing air as the heat exchange medium in direct contact with the heat load. As it is already in liquid form, it is very

efficiently matched with the already described closed water loop serving the secondary heat rejection plant. The interface between the primary and secondary fluid loops comes in the form of a Cooling Distribution Unit (CDU).

The primary fluid system is filled once when the cooling plant is commissioned and replenished once or twice a year for maintenance or if a leak occurs. It is usually filled with a mixture of water and chemical additives to reduce microbial buildup and ensure increased performance. When this mixture is changed, it is either treated onsite or at a local water facility, meaning there is practically no water consumption (see Myth 3).



[MYTH 2]


“All data centers cool with water”

Which data centers? Built when and where? Cooled how?

The differences between cooling technologies and their associated water consumption figures are not marginal. [Google’s own facility-level disclosures for 2024](#), a useful example of the transparency this debate needs more of, show one Texas data center consuming about 10,000 gallons in a year, roughly two months of a single household’s use, while its largest water consumer, in Iowa, used around a billion gallons. Same company, same year, a difference of five orders of magnitude, determined almost entirely by cooling architecture and climate.

One caveat: modern data center cooling designs are not representative of the whole. Data centers have been around for decades, and the installed base is huge. In many countries, evaporative cooling remains widespread across the existing global fleet, and [Microsoft](#)

[has noted](#) that its current facilities will remain a mix of air-cooled and water-cooled systems even as its post-August 2024 designs for new facilities eliminate evaporative water use. When someone says, “data centers don’t really use water” or “data centers devour water,” the right response is to ask: which data centers? Built when and where? Cooled how? Both statements are true of some facilities and false of others.



Grouped sources for this section can be found on [pages 29-30](#), sections [05] and [07]

[MYTH 3]

“Water used
by a data center
is gone forever”

10 MYTHS ABOUT
AI DATA CENTERS
AND WATER

[MYTH 3]

“Water used by a data center is gone forever”

Withdrawal is not the same as consumption.

The single most conflated distinction in this debate is between withdrawal, meaning water taken from a source, discharge, meaning water used and then returned to a waterway or treatment facility, and consumption, meaning water lost to evaporation. The source of the withdrawal is also important to note. A facility that withdraws river water, runs it through heat exchangers, and returns it at a slightly higher temperature has a very different impact from one that draws from aquifers and evaporates the same volume into the sky. Colocation provider [Equinix, for instance, reported](#) withdrawing 1.4 billion gallons globally in 2024 while consuming 1.2 billion. Many press accounts of data center water use quote withdrawal figures as if they were consumption, or vice versa, without noticing the difference.

The distinction refines the problem rather than dissolving it. Returned does not mean impact-free; thermal

discharge from water that has been through the data center cooling process affects river ecosystems, and the timing of withdrawals matters during drought. Water can also be contaminated during use in a data center, requiring proper treatment before it is discharged back into the watershed. On the other hand, consumed water genuinely is lost from the local source. That water eventually falls as rain somewhere on Earth, but that is of no comfort to the aquifer it was taken from, especially if that aquifer is also used for watering crops by several farms, or for drinking water in a nearby city. The lesson is not that withdrawal does not count. It is that a report which does not tell you which figure it is quoting has not told you the full story.

What the numbers say

WITHDRAWAL	DISCHARGE	CONSUMPTION
taken	returned	lost
from a source	after use	to evaporation

THE CAVEAT

Returned does not mean impact-free. A report which does not tell you which figure it is quoting has not told you the full story.

Grouped sources for this section can be found on [page 29](#), section [05]

[MYTH 4]

“AI’s water footprint
is mostly at
the data center”

10 MYTHS ABOUT
AI DATA CENTERS
AND WATER

[MYTH 4]

“AI’s water footprint is mostly at the data center”

Zero water on-site can still mean water demand upstream and out of sight.

Here is a myth that cuts against the industry, and it deserves far more attention than it gets. [The 2024 US Data Center Energy Usage Report](#) from Lawrence Berkeley National Laboratory estimated that American data centers consumed about 17 billion gallons of water directly through cooling in 2023, and roughly 211 billion gallons indirectly, through the water consumed generating their electricity. On those figures, more than 90% of the sector’s water footprint never appears in an operator’s on-site metering, because thermoelectric generation, including nuclear, is among the most water-intensive activities in the economy, [accounting](#) for around 34-38% of US freshwater withdrawals, and hydropower reservoirs add further evaporative losses under some accounting methods.

This means a data center that consumes zero water on-site, while

drawing power from a water-hungry grid, may have simply moved its water demand upstream and out of sight.

The power generation figure comes with a multitude of caveats and is genuinely uncertain. The LBNL estimate has been credibly challenged as potentially overstating indirect water consumption by a factor of two or more, depending on assumptions about grid mix, and it varies enormously by location. A facility running on a wind and/or solar-heavy grid has a dramatically smaller upstream footprint than one on conventional thermal generation, and data centers can rarely choose where the electrons in their grid originate. But the overall thrust of the report is robust across every serious analysis: for most facilities, electricity is where the majority of the water use is. Any assessment that stops at the facility fence is measuring the smaller number.

What the numbers say

US DATA CENTERS / 2023 ESTIMATE

17B gal

direct water through on-site cooling

211B gal

indirect water through electricity generation

>90%

outside on-site metering

THE CAVEAT

The indirect estimate is uncertain and varies by grid mix and location. Any assessment that stops at the facility fence is measuring the smaller number.

Grouped sources for this section can be found on [page 28](#), sections [02] and [03]

[MYTH 5]

“Any water use
is bad water use”

10 MYTHS ABOUT
AI DATA CENTERS
AND WATER

[MYTH 5]

“Any water use is bad water use”

Water stress is local in a way carbon is not.

If Myth 4 establishes that electricity has a hidden water cost, this is its corollary: sometimes using water on-site is the environmentally responsible choice. Evaporative cooling exists because it is significantly more energy-efficient than mechanical alternatives in many climates. Microsoft acknowledged this trade-off explicitly when announcing its zero-water designs, noting that moving away from evaporation increases a facility’s power usage effectiveness (PUE), the ratio between power used for support systems like cooling and power used for IT. There are other ways of reducing PUE, but water is one of the most effective.

Given the water-intensive and carbon-intensive nature of some forms of power generation, in a genuinely water-abundant region a facility that evaporates water on-site but burns meaningfully less electricity can have a lower total water and carbon footprint than a “water-free” neighbor pulling harder on the grid. Water stress is local in a way carbon is not.

A liter evaporated beside a rain-fed river system and a liter pumped from a depleting desert aquifer are not ecologically equivalent, and treating all liters as equal is how bad siting decisions get presented as good ESG.

This argument applies only where water abundance is real and verified against local hydrology, meaning river flows, aquifer recharge, seasonal drought risk, and, of course, the effects of a changing climate on the watershed. It should never act as a license to site evaporative cooling in stressed basins because the energy math works on paper. The right question is not “does it use water?” but “what is the full water-energy balance for this site, on this grid, and in this watershed?” The answer should err on the side of caution and be verifiable by parties beyond the operator alone.

What the numbers say

EVAPORATIVE COOLING

+ water
- energy

ZERO-WATER COOLING

- water
+ energy

THE BEST OPTION

depends on the site

THE CAVEAT

This applies only where water abundance is real and verified. It is not a license to use evaporative cooling in stressed basins just because the energy math works on paper.

Grouped sources for this section can be found on [page 30](#), section [07]

[MYTH 6]

“Data centers are
uniquely enormous
water users”

10 MYTHS ABOUT
AI DATA CENTERS
AND WATER

[MYTH 6]

“Data centers are uniquely enormous water users”


Small nationally.
Significant locally.

In most countries, at the level of national accounting, data centers barely register. For the US, the LBNL figure mentioned in Myth 4 of roughly 17 billion gallons of direct annual consumption works out to about 0.3% of the US public water supply, per analysis by [Ren and Luers via the IEEE Spectrum](#). For comparison, the Golf Course Superintendents Association of America's own [survey data](#) puts US golf course irrigation at roughly 1.5 to 2 billion gallons per day. Golf uses approximately thirty times more water than every American data center combined. Agriculture dwarfs both: the [UN's World Water Development Report](#) puts farming at roughly 70% of global freshwater withdrawals. A typical large data center consumes on the same order as a single golf course; one facility in Virginia used

about 173 million gallons in a year, on the order of a single golf course's irrigation.

Three caveats on these figures are worth bearing in mind before anyone in the industry takes too much comfort from them.

First, the comparison figures are softer than they look. The 70% agriculture number, though endorsed by UN reporting, has been [challenged in academic literature](#) as poorly evidenced, with plausible values anywhere from 45% to 90%. [Facility-level data center consumption](#) spans five orders of magnitude (see Myth 2), so “typical” conceals more than it reveals. The data is still too patchy to use conveniently chosen endpoints.




[MYTH 6]

“Data centers are uniquely enormous water users”

Golf uses more is context, not exoneration.

Second, and more fundamentally, national figures obscure local impact. Water is consumed from local watersheds, not from a national reservoir. A single hyperscale facility drawing millions of gallons a day, arriving suddenly in one small municipality, can matter more locally than a thousand diffuse legacy uses spread across a continent. This is precisely why the concern exists despite the aggregate numbers. Data centers attract scrutiny not because they are large water users nationally but because they are new, concentrated, and visible, and those are rational things for a community to scrutinize.

Third, “golf uses more” is context, not exoneration. Golf courses are themselves widely criticized for water use, so matching a criticized benchmark is a low bar. And on current growth projections, particularly for AI data centers, the comparison erodes: [LBNL projects](#) direct data center consumption could double or quadruple by 2028. The honest framing is that data centers are currently a small national water user with genuinely large local footprints and a steep growth curve. All three clauses matter.



[MYTH 7]

“More compute
always means
more water”

10 MYTHS ABOUT
AI DATA CENTERS
AND WATER

[MYTH 7]

“More compute always means more water”

More compute does not automatically mean more water.

The assumption that more AI means more heat, which, in turn, means more evaporation runs into an engineering reality: extreme rack densities are exactly what evaporative systems were not designed for. The newest generation of AI facilities are increasingly built around closed-loop, direct-to-chip, and immersion cooling, precisely because the thermal loads demand it. [Microsoft's post-2024 designs](#) eliminate evaporative cooling entirely; its fleet-wide water usage effectiveness improved from 0.49 to 0.30 liters per kilowatt-hour between 2021 and 2024. [Google reported](#) a 33-fold reduction in energy per median Gemini prompt in a single year, with water falling correspondingly. Per unit of compute, the trajectory of water intensity is sharply downward, and the densest AI workloads are, perhaps counterintuitively, often served by the least evaporative cooling architectures.

The important caveat: intensity is not volume. Efficiency per prompt can fall while total consumption rises if growth outruns the gains, the classic Jevons dynamic, and the projections say exactly that. [LBNL expects](#) US direct consumption to grow two-to-four-fold by 2028, even as new builds get more efficient. The fleet is expanding faster than it is improving. The myth worth busting is the automatic link between compute growth and water growth, not the possibility that sector totals rise. On current evidence, they will.

What the numbers say

PER UNIT ↓

0.49 → 0.30 L/kWh

Microsoft fleet-wide WUE
2021–2024

TOTAL VOLUME ↑

2–4× higher

projected US direct water
consumption by 2028

THE CAVEAT

Intensity is not volume. Total consumption can rise when growth outpaces efficiency gains.

Grouped sources for this section can be found on [pages 28–30](#), sections [01], [02], and [07]

[MYTH 8]

“Data centers take drinking water from communities”

10 MYTHS ABOUT
AI DATA CENTERS
AND WATER

[MYTH 8]

“Data centers take drinking water from communities”

A documented case matters. It does not define the whole industry.

Stated as a general rule about the industry, this is false; the majority of facilities operate without measurable effect on municipal supply. But this myth demands the most careful handling in the list, because in specific, documented cases community concerns have been well founded, and acknowledging that plainly is essential to the credibility of everything else in this paper.

The reference case in this debate is European. In North Holland, in the Netherlands, a hyperscale campus at Middenmeer [was reported in 2022](#) to have consumed 84 million liters of drinking water in a year, against a projected 12 to 20 million, with the figures emerging during a declared national drought and after sustained local opposition to the site’s water use. The operator attributed much of the difference to construction-phase water and noted that a substantial volume was returned to the local supply, but the gap between estimate and outcome became a defining episode in the European debate over consumption forecasting.

What a case like Middenmeer does not establish is that it is

representative of the industry today. It reflects a facility sited and permitted in an earlier era, before water became a first-order design constraint, a permitting priority, and a boardroom-level consideration. The industry’s response since then has been substantive: new builds increasingly arrive with closed-loop designs, reclaimed-water agreements, and materially stronger disclosure commitments attached to their permits, and the largest operators have led much of that shift in design standards.

The word “increasingly” in the last paragraph is doing honest work, but it is a trend claim, not a guarantee, and it is exactly the kind of claim that should be verified permit by permit and site by site rather than taken on assurance. The correct lesson for communities, operators, and planning authorities is not that the problem is imaginary. It is that the problem is preventable, the engineering to prevent it already exists, and early, transparent engagement is now the established best practice for deploying it.

What the numbers say	PROJECTED	REPORTED	DIFFERENCE
	12–20M L	84M L	≈4–7x
	drinking water / year	Middenmeer / 2021	projected
THE CAVEAT			
Middenmeer is a documented case, not proof that it represents the industry today. Improvements must still be verified permit by permit and site by site.			

Grouped sources for this section can be found on [page 29](#), section [06]

[MYTH 9]

“Water use is just the
cost of progress, and
nothing can be done”

10 MYTHS ABOUT
AI DATA CENTERS
AND WATER

[MYTH 9]

“Water use is just the cost of progress, and nothing can be done”

The solutions already exist. The challenge is making them common practice.

Almost everything can be done, and much of it is already commercial rather than experimental. Closed-loop and direct-to-chip cooling eliminate evaporative consumption for new builds. Microsoft’s zero-water designs come online from 2027, saving a stated 125 million liters of water per facility per year. Reclaimed and non-potable water can replace drinking water at facilities that do use some form of evaporative cooling. Warmer coolant temperatures, enabled by the tolerances of modern chip families such as NVIDIA’s Rubin generation, recover the energy penalty of choosing dry cooling over wet. Water usage effectiveness (WUE) gives regulators and communities a metric to write into permits. And the three largest hyperscalers, Microsoft, Google, and Amazon, have all committed to being “water positive” by 2030, replenishing more than they consume.

One caveat has already been mentioned: global figures should not eclipse local impact. Replenishment commitments are global accounting, and restoring water to one basin does not refill the aquifer beside a specific facility. Critics, including Ren, have made exactly this point, and it is valid. Best practice and common practice also remain some distance apart: the solutions above are unevenly adopted, mostly voluntary, and concentrated among the best-resourced operators. The gap between what the technology permits and what the median facility does is itself part of the problem, and closing it is a matter of permitting and procurement standards, not invention.

What the numbers say

ROLLOUT

2027

Microsoft zero-water designs

→

WATER SAVING

125M L

per facility / year

→

WATER-POSITIVE TARGET

2030

Microsoft, Google & Amazon

THE CAVEAT

Global replenishment does not refill the same local basin. Best practice remains unevenly adopted and mostly voluntary.

Grouped sources for this section can be found on [page 30](#), section [07]

[MYTH 10]

“The numbers
can be taken
at face value”

10 MYTHS ABOUT
AI DATA CENTERS
AND WATER

[MYTH 10]

“The numbers can be taken at face value”

Scrutiny is not cynicism. It is the rational posture.

Not yet, and this paper’s own sourcing shows why. Not every major operator publishes aggregate water figures. Google’s five-drops-per-prompt measurement counts only on-site water, a boundary choice that excludes the majority of the footprint identified in Myth 4; to Google’s credit, that boundary is disclosed in the report itself, but it is one the media headlines routinely omit. There is no mandatory water-disclosure requirement for data centers in the United States at federal level, and WUE reporting, where it exists, uses inconsistent boundaries that make operators difficult to compare.

Scrutiny of the numbers is therefore not cynicism; it is the rational posture until reporting is standardized, mandatory, and independently verifiable. The European Union has moved further than most: under the recast [Energy Efficiency Directive](#), data centers above 500kW of installed IT power have been [required](#) to report annual KPIs,

including water usage effectiveness, to a European database since 2024, and the European Commission has adopted the first phase of an EU-wide sustainability rating scheme. Public disclosure, however, remains aggregated at country and EU level rather than facility level, except where national implementation goes further, as in the Netherlands, where [legislation](#) makes individual operators’ energy and water consumption public. Despite room for improvement, legislation of this kind is a straightforward way to improve transparency and, ultimately, public confidence in the industry.

That scrutiny should be applied symmetrically to the viral bottle-of-water figure as much as to the corporate five drops. Operators hold the meters, which is why leadership on disclosure sits naturally with them, and why the operators that proactively standardize ahead of legislative requirements will set the terms of trust.

What the numbers say

UNITED STATES	EUROPEAN UNION	NETHERLANDS
no federal	>500 kW	facility-level
mandatory water disclosure	annual KPI reporting required since 2024	energy and water data public

THE CAVEAT

Reporting boundaries remain inconsistent. Even where disclosure is required, public data is often aggregated rather than facility-level.

Grouped sources for this section can be found on [page 30](#), section [07]

The real questions to ask

Every number in this debate should be read as provisional: an estimate with a boundary, not a fact with an audit.

The question that matters was never just “does AI use water?”
The questions that matter are:

- **What?**
Is the cooling methodology evaporative, air, or liquid?
- **Which?**
Is the water potable or reclaimed, evaporated or returned / discharged?
- **Where?**
An abundant watershed or a stressed one?
What is the water intensity of the regional grid mix?
- **How much?**
- **When?**
Including time of day, drought months or years, or just overall averages?
- **Who checks?**
Anyone outside the company, or is data provided via report taken at face value?

On today's evidence, AI data centers are a small national water user with a real and rising footprint, genuinely severe in a handful of badly sited legacy cases, and increasingly solvable by planning and engineering that already exist. What is not yet solved is the accounting. Standardized, facility-level, independently verifiable disclosure of withdrawals, discharge, and consumption, direct and indirect, would cost the industry relatively little and would do more to build trust with the communities hosting it than any single commitment. Until that exists, every number in this debate, including the reassuring ones, should be read as provisional: an estimate with a boundary, not a fact with an audit.

Sources and boundaries (for editors and fact-checkers)

[01] Per-query figures:

Li, Yang, Islam, and Ren, Making AI Less “Thirsty” (arXiv:2304.03271, first posted April 2023, <https://arxiv.org/abs/2304.03271>; later published in Communications of the ACM, Vol. 68 No. 7, 2025, <https://cacm.acm.org/sustainability-and-computing/making-ai-less-thirsty/>), origin of the 500ml per 10 to 50 GPT-3 responses estimate; The Washington Post (September 18, 2024, with Ren), approximately 519ml per 100-word GPT-4 email (<https://www.washingtonpost.com/technology/2024/09/18/energy-ai-use-electricity-water-data-centers/>); Ren’s subsequent revision to approximately 15ml per GPT-4 prompt, approximately 5ml on-site (documented correspondence published May 2026, <https://blog.andymasley.com/p/i-might-have-found-the-specific-way>; corroborated by Forbes, July 2026, <https://www.forbes.com/sites/robertszczerba/2026/07/21/how-much-water-does-ai-use-the-58-billion-risk/>); Google, Measuring the environmental impact of delivering AI at Google Scale (August 2025, arXiv:2508.15734, <https://arxiv.org/abs/2508.15734>; companion post at <https://cloud.google.com/blog/products/infrastructure/measuring-the-environmental-impact-of-ai-inference>), 0.26ml median Gemini prompt, on-site boundary only. These estimates use different boundaries and are not directly comparable; that is part of the point.

[02] US totals and projections:

Lawrence Berkeley National Laboratory, 2024 United States Data Center Energy Usage Report (LBNL-2001637, December 2024, https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report_1.pdf), 17 billion gallons direct consumption (2023), approximately 211 billion gallons indirect via electricity, direct consumption projected to reach roughly 34 to 73 billion gallons by 2028 depending on scenario. The indirect estimate has been publicly challenged as potentially high by approximately 2x, partly on the treatment of hydropower reservoir evaporation (Construction Physics, September 2025, <https://www.construction-physics.com/p/i-was-wrong-about-data-center-water>); treat as an upper-bound-leaning estimate with a robust directional conclusion.

[03] Sector shares:

UN World Water Development Report 2024, agriculture approximately 70% of global freshwater withdrawals (<https://www.unesco.org/reports/wwdr/en/2024/s>); note this canonical figure has been academically contested, with a plausible range of 45% to 90% (Puy et al., PNAS Nexus, November 2025, <https://academic.oup.com/pnasnexus/article/4/11/pgaf323/8320029>); USGS, thermoelectric power approximately 40% of US freshwater withdrawals (34% to 38% depending on survey year, <https://pubs.usgs.gov/publication/cir1405>);

Ren and Luers, The Real Story on AI Water Usage at Data Centers (IEEE Spectrum, 2025, <https://spectrum.ieee.org/ai-water-usage>), US data centers approximately 0.3% of contiguous-US public water supply, as cited in AGU Advances (<https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2025AV002140>).

[04] Golf comparison:

Golf Course Superintendents Association of America survey data, approximately 1.68 million acre-feet applied in 2020 (<https://www.gcsaa.org/who-we-are/media/news/2022/07/26/golf-courses-reduce-water-usage-by-29-percent-according-to-national-survey>; full report at https://www.gcsaa.org/docs/default-source/environment/22_waterreport_web.pdf). Both the golf and data center figures are estimates with different methodologies; the approximately 30x ratio is robust to that uncertainty, the precise multiple is not.

[05] Facility-level variance:

Google 2025 Environmental Report and facility disclosures, Council Bluffs, Iowa, approximately 1 billion gallons (2024) versus Pflugerville, Texas, approximately 10,000 gallons (air-cooled); Leesburg, Virginia, 173.2 million gallons (2023), per the Google 2024 Environmental Report (<https://www.gstatic.com/gumdrop/sustainability/google-2024-environmental-report.pdf>; figure as reported by Quartz, February 2025, <https://qz.com/google-data-center-consume-water-energy-climate-change-1851760734>). Equinix 2024 disclosures, 1.4 billion gallons withdrawn, 1.2 billion consumed (<https://www.equinix.com/resources/data-sheets/sustainability-data-summary>).

[06] Community cases:

Middenmeer, North Holland, Netherlands, approximately 84 million liters of drinking water consumed in 2021 versus a projected 12 to 20 million, reported by Noordhollands Dagblad (August 2022, https://www.noordhollandsdagblad.nl/cnt/dmf20220810_68483787) during a declared national drought; the operator attributed a substantial share to construction-phase use and stated approximately 36 million liters were discharged back into the regional supply (Datacenter Dynamics, <https://www.datacenterdynamics.com/en/news/drought-stricken-holland-discovers-microsoft-data-center-slurped-84m-liters-of-drinking-water-last-year/>).

[07] Cooling technology and commitments:

Microsoft Cloud Blog (December 9, 2024, <https://www.microsoft.com/en-us/microsoft-cloud/blog/2024/12/09/sustainable-by-design-next-generation-datacenters-consume-zero-water-for-cooling/>), zero-water-evaporated designs from August 2024, pilots in Phoenix, AZ, and Mount Pleasant, WI, online from late 2027, approximately 125 million liters per year saved per facility, fleet WUE 0.49 to 0.30 L/kWh (2021 to 2024), acknowledged PUE trade-off; water-positive-by-2030 commitments: Microsoft (2020, <https://blogs.microsoft.com/blog/2020/09/21/microsoft-will-replenish-more-water-than-it-consumes-by-2030/>), AWS (2022, <https://www.datacenterdynamics.com/en/news/aws-pledges-to-be-water-positive-by-2030/>), and Google (<https://www.computerweekly.com/news/252506608/Google-joins-Microsoft-and-Facebook-by-pledging-to-replenish-more-water-than-it-consumes-by-2030>). Regulation: EU Energy Efficiency Directive (recast, 2023/1791, <https://eur-lex.europa.eu/eli/dir/2023/1791/oj>) and Delegated Regulation 2024/1364 (https://eur-lex.europa.eu/eli/reg_del/2024/1364/oj/eng), mandatory annual KPI reporting, including WUE, for data centers above 500kW of installed IT power; first reports to the European database by September 2024, then by May 15 annually; public disclosure aggregated at EU and member-state levels; facility-level transparency in the Netherlands under national implementation (NL Times, July 2026, <https://nltimes.nl/2026/07/14/microsoft-data-center-uses-1-dutch-electricity>).

All figures are estimates carrying the boundaries stated by their sources; where sources conflict, the piece states the range. No figure in this paper originates from Submer Group or its subsidiaries.

Connected
intelligence for
AI infrastructure

→ [Learn more at our blog](#)



Submer Group