

The Cost of AI, Seen Through the Lens of Water

A Studio Wā Living Memo

Water, wealth, and the price of the AI buildout.

How, not whether

This memo is not about whether you should use AI. That question is mostly already answered, and not by you. If you do anything in the digital world, you are using AI. It is in the search bar, the map, the phone keyboard, the spam filter, the bank's fraud check, the photo app, and the feed that chooses what you see.¹ Choosing not to subscribe to ChatGPT or Claude opts you out of a few visible tools. It does not opt you out of AI. Treating personal abstinence as the whole of the ethical question is a comfort, and a small one.

The question worth your attention is not whether, but how: how to make an informed decision about the tools you can actually choose, with the right data in front of you.

Studio Wā holds this as a human problem first. The future worth building is a sustainable, equitable, and more loving one for everyone, and finding the way there is work humanity has to do together. We center Indigenous peoples inside that larger frame for two reasons that run toward each other. The harm lands on Indigenous communities first and hardest, which makes the stakes most pressing there. And the wisdom most needed to find a way through, intergenerational knowledge tested across centuries of upheaval, has survived most intact in those same communities. Most cultures once carried some version of it. For many it has been interrupted rather than lost. Centering Indigenous peoples is not a narrowing of the human project but a way back into it, and for this studio it is also the taproot, the ground the rest grows from.

The harm, multiplied

AI rarely invents a new way to harm Indigenous peoples. It takes the old ways, the ones colonization has run for centuries, and speeds them up, automates them, and runs them all at once. Data gets taken without consent, now at the scale of entire training sets, and at the same time Indigenous knowledge is mostly kept out of how these systems are built, so a benefit that could be real is never

allowed to form. Languages get modeled, overwhelmingly by outsiders rather than the communities themselves, and risk being frozen in place by people who do not speak them.² Cultural knowledge, design, and art get scraped and regenerated with no credit and no return. Land and water get seized for infrastructure. Surveillance gets cheaper and more total. None of these is new. What is new is the speed, the scale, and the fact that they arrive together. That is what a threat multiplier is. It does not add one danger to the pile. It raises the level of every danger already in the room.

A multiplier is hard to write about, because the honest version is a list of everything at once, and a list of everything is the same as no focus at all. So this memo takes one front, water, as its lens. Not because water ranks first, and not because the others matter less, but because water is the one harm that cannot be abstracted away, and because in a Hawaiian sense water is life and waiwai is wealth, which makes it the right place to begin.

You can argue about a language model. You cannot argue with an empty aquifer.

Water is life

Water is the one resource with no substitute and no patience, and no community can wait on a well. As AI data centers get built, two things are true at once and they pull against each other. One honest option is to stop building them. Until that happens, they go somewhere, and wherever they go, they have to drink. Power, land, heat, money, and consent all gather around that single demand for water. As the studio leans further into its Indigenous taproot, water comes to look less like one issue among many and more like the place the others meet. It is the one wealth every living thing shares, and the medium life itself began in.

It may even be the chain that links them. Take the water and you have already taken the food it grows, the health it carries, the practices it makes possible, and in many cases the ability to stay on the land at all. That is one way to see it, not a proven law, but it is the way that matters most to people whose relationship to place is rooted in water first, or in lands whose abundance depends on it.

The numbers, in context

Someone will tell you that one AI query used a bottle of water. Someone else will tell you it used less than a dishwasher cycle, so relax. Both could be technically defensible, but both distract from the context, which matters more.

The per-query number really is tiny. A short request runs on a fraction of a watt-hour and a fraction of a milliliter of water. Even the high academic estimate, around half a liter for a hundred words of generated text,³ means one load of laundry equals dozens of prompts. So no, the picture of a river draining every time you open a chat window is not true. But that picture is wrong only from the vantage

of a single person's use, which is the same vantage the dishwasher comparison takes. Both are accurate for one person on one afternoon, and both tell you nothing about the system that afternoon belongs to: who else is using it, and what else is moving through the AI ecosystem that same day. That individual framing is the fossil fuel industry's personal-carbon-footprint move, relocating a structural problem onto private conscience. Clarity arrives only when we move from the individual to the aggregate.

The unit that matters is the system. Global data centers consumed an estimated 448 terawatt-hours of electricity in 2025, which on its own would rank as the world's eleventh-largest national consumer.⁴ A UN report released this month projects that AI data centers could draw water and power on the scale of 1.3 billion people's use by 2030.⁵ A single midsize facility can consume about 110 million gallons of water a year, roughly a thousand households. A large one can reach five million gallons a day.⁶

Three things separate this from ordinary unsustainable living. The demand is additive: it stacks on top of what we already use, grows faster than any other category, and replaces nothing. It is also concentrated and sited without consent. A household's use spreads thin across a grid, while a data center's use lands on one watershed and one community, often rural, water-stressed, or Indigenous, chosen without the community's say. And the benefit and the burden run in opposite directions. The profit concentrates in a few firms. The water, heat, noise, and grid strain stay local. Roughly two-thirds of new data centers in the United States are going up in places already short of water,⁷ and Native land is increasingly where developers go looking, from the Arizona desert to the Plains to Virginia.

The honest stance is neither guilt nor permission. Your individual use is not the harm, and abstaining from it is not the remedy. The questions worth carrying are relational:

- What does this take, and from whom?
- What returns?
- Who decided?

Waiwai, our collective wealth

The Hawaiian word for wealth is waiwai. Its root, wai, is water. Wealth was the water itself: what grows the food, what sustains the people, what makes abundance possible at all, the same truth the cry that water is life carries now.

The companies building this infrastructure measure wealth a different way, in cash, market value, and digital currency, and they gather it by the logic they apply to everything else: take as much as possible, hold it, release as little as you can. Applied to money, we call that hoarding. Applied to water, it has no comfortable name, but it is the same act. The water is drawn down and locked inside the machine's demand, and what flows back to the community is not water and not abundance. At best it is a smaller

amount of money. Real wealth gets converted into the abstract kind and carried off somewhere else. That is a wealth crisis in the oldest sense, even when the financial statements read as growth.

A water-sovereignty question

Water on Indigenous land has a long history of being allocated by everyone except the people who live with it. Data centers arrive inside that history, not outside it. They come with the familiar promises, jobs and local investment, and organizers in Indian Country say those mostly fail to show, while the water draw, the rate hikes, and the risk of displacement do not.⁸ In parts of Arizona, families live without reliable power while data centers rise nearby.⁹

This is why the studio keeps water next to data sovereignty rather than in a separate environmental file. The legal handle is the same one that governs Indigenous data: free, prior, and informed consent, and the right under UNDRIP Article 31 of peoples to control what is done with and on their territories. A watershed is territory. The question of who may build a five-million-gallon-a-day machine on it is a sovereignty question before it is an engineering one.

Can AI help with the harm it creates?

Partly, and the honest answer carries a trap. Real applications exist. Machine-learning cooling control cut one large operator's cooling energy by around 40 percent.¹⁰ Forecasting helps integrate renewable power. AI supports climate modeling, conservation monitoring, and language revitalization, the last of these proven by Te Hiku Media in te reo Māori on Māori terms.¹¹

The trap is Jevons' paradox: efficiency gains have historically raised total consumption rather than lowering it, because the saving gets reinvested in more use.¹² Efficiency is not a value. It is a capacity. It can be pointed at restoration or at extraction with equal ease. The Indigenous frame asks the question efficiency cannot: directed toward what, in reciprocity with whom, accountable across how many generations.

What this memo leaves open

Water is one front, taken as a lens because it cannot be abstracted. The multiplier reaches further, and the studio means to keep writing toward the rest. The fuller map, AI as a threat multiplier across all these fronts, is its own memo to come.

One thread needs singling out, because it is where a water memo keeps wanting to wander. Beyond the harm of extraction lies the cost of absence: Indigenous knowledge is largely kept out of how these models are built and reason, so a benefit that could be real never forms. That conversation is already alive on the ground, and it holds several positions at once rather than choosing between them: that intergenerational knowledge and cultural property not be used for training without consent, that the

community govern any use that does happen, and that outright refusal stay fully valid in parallel. Running alongside is a newer question, whether a decision-making model could be built on Indigenous terms rather than scraped from them. The studio's stance is that this is the work most worth doing, and that it earns its own memo rather than a paragraph here.

The rest of the map, named so it is not mistaken for the whole: the re-scraping of data already taken without consent, the generative copying of Indigenous art and story, language sovereignty beyond the few community-governed models that exist, the surveillance that cheap AI makes total, and the human labor, often paid least, that cleans the data these systems learn from. Each is its own memo. Naming them here is how we say that water is where we start, not where the harm ends.

Staying current

No single document stays true in this field for long, so the more useful thing than any fact above is a habit. A suggested rhythm, roughly fifteen minutes at each interval.

Every two weeks, the news layer. Scan two Indigenous-led outlets for what is being built and contested right now. Grist's Indigenous desk (grist.org/indigenous) reports steadily on AI, land, water, and knowledge extraction in Indian Country. ICT News (ictnews.org) covers Indian Country directly, including data-center siting on tribal land.

Once a month, the rights-and-sovereignty layer. Check where the protective frameworks and policy are moving. Cultural Survival (culturalsurvival.org) follows Indigenous rights and technology, including its AI coverage and quarterly magazine. The Global Indigenous Data Alliance (gida-global.org) stewards the CARE Principles, and Te Mana Raraunga in Aotearoa and the US Indigenous Data Sovereignty Network carry the regional developments.

Once a quarter, the depth layer. Read one longer thing to keep the analysis sharp rather than reactive. For the constructive side, Abundant Intelligences (abundant-intelligences.net) and Te Hiku Media (tehiku.nz) on Indigenous-governed AI. For the cost side, the water research of Shaolei Ren at UC Riverside and the periodic UN and IEA reports on data-center water and energy.

A practical note on the water numbers specifically: treat any single figure as a snapshot, note its date, and re-check it against the quarterly reports before repeating it. The studio's own figures here are current as of June 2026.

About Studio Wā living memos. A living memo is the studio's account of a topic as of the date it carries. Living means it will be revised as our learning grows, from written sources and from conversation, which the ground here requires.

Why this is Studio Wā work. Studio Wā is a Hawaiian-rooted studio concerned with time: how cultures build it, hold it, and keep it their own. AI can look like a departure from that subject. It is the opposite. At its core AI is a temporal force. It collapses the time a task takes, removes the human from

the doing, and resets what a unit of time is worth, which drives the same colonization of time the studio exists to name. It also scrapes and reshapes the cultural and temporal knowledge peoples have carried across generations, and Hawaiian timekeeping is among the most fully developed temporal infrastructures any society has made. Asking how that knowledge is held, and who AI serves and at what cost, is not a digression from the studio's work on time. It is the present-day front of it.

References and further reading

Verified June 2026. The numbered citations above are collected here in full, followed by two readings for going further. All links are live in the Word and PDF exports.

- Li, P. et al. *Making AI Less “Thirsty”: Uncovering and Addressing the Secret Water Footprint of AI Models*. UC Riverside, 2023. arxiv.org/abs/2304.03271
- *Energy, Water Use, and Pollution of AI and Data Centers Rival Most Countries*. PBS NewsHour. pbs.org/newshour
- *AI and Global Water Resources*. UN report coverage, June 2026, via Time. time.com
- *Data Centers and Water Consumption*. Environmental and Energy Study Institute. eesi.org
- *Most New U.S. AI Data Centers Are Being Built in Drought Zones*. Tom's Hardware. tomshardware.com
- *In Indian Country, Data Centers Come with a Familiar Threat of Colonialism*. ICT News. ictnews.org
- *Amid Arizona's Data Center Boom, Many Native Americans Live Without Power*. Washington Post. washingtonpost.com
- *AI's Climate-Positive Applications and Jevons' Paradox*. ACM FAccT 2025. dl.acm.org
- *Data Sovereignty and the Kaitiakitanga License*. Te Hiku Media. tehiku.nz
- *Introducing Speech-to-Text, Text-to-Speech, and More for 1,100+ Languages (Massively Multilingual Speech)*. Meta AI, 2023. ai.meta.com
- *This Is Why Half of the World's Languages Are Endangered*. World Economic Forum, 2022. weforum.org
- *Key Findings About How Americans View Artificial Intelligence*. Pew Research Center, 2026. pewresearch.org

To go further:

- Carroll, S. R. et al. *The CARE Principles for Indigenous Data Governance*. Data Science Journal, 2020. The companion framework for who controls Indigenous data. doi.org/10.5334/dsj-2020-043

- Lewis, J. E., Arista, N., Pechawis, A., and Kite, S. *Making Kin with the Machines*. Journal of Design and Science, MIT Press, 2018. The founding essay of the Indigenous AI conversation. jods.mitpress.mit.edu
-

1. AI and machine learning now run quietly inside everyday digital services: search ranking, recommendation feeds, spam filtering, fraud detection, computational photography, and predictive text. On its growing presence in daily life, see Pew Research Center, 2026. pewresearch.org↵
2. Of roughly 7,000 living languages, about 4,000 are spoken by Indigenous peoples, and an estimated 97 percent are “digitally disadvantaged.” The large multilingual systems are built by corporations rather than communities: Meta’s Massively Multilingual Speech project alone provides models for over 1,100 languages and data for nearly 4,000, assembled by the company. Community-governed efforts such as Te Hiku Media’s te reo Māori models remain rare. Meta AI, 2023, ai.meta.com; World Economic Forum, 2022, weforum.org.↵
3. Pengfei Li et al., *Making AI Less “Thirsty,”* UC Riverside, 2023. arxiv.org/abs/2304.03271↵
4. *Energy, water use, and pollution of AI and data centers*, PBS NewsHour. pbs.org/newshour↵
5. AI and global water resources, UN report, June 2026, via Time. time.com↵
6. *Data Centers and Water Consumption*, Environmental and Energy Study Institute. eesi.org↵
7. *Most new U.S. AI data centers are being built in drought zones*, Tom’s Hardware. tomshardware.com↵
8. *In Indian Country, data centers come with a familiar threat of colonialism*, ICT News. ictnews.org↵
9. *Amid Arizona’s data center boom, many Native Americans live without power*, Washington Post. washingtonpost.com↵
10. Google DeepMind reported an approximately 40 percent reduction in data-center cooling energy from machine-learning control, 2016.↵
11. *Data Sovereignty and the Kaitiakitanga License*, Te Hiku Media. tehiku.nz↵
12. *AI’s climate-positive applications and Jevons’ paradox*, ACM FAccT 2025. dl.acm.org↵