OLDS RESEARCH

Scare Gallons vs. Real Gallons

What the water fight gets wrong about AI data centers

August 2026

Executive Summary

New AI data centers are being blocked in hundreds of U.S. communities on the claim that they will drink municipal water dry. The claim does not survive contact with national data or with how those buildings are actually cooled.

On the U.S. Geological Survey's national survey of water usage, the big draws are thermoelectric power (~133 billion gallons a day), irrigation (~118), public supply (~39), and industry (~15).

Lawrence Berkeley National Laboratory's recent assessment puts direct data-center water use at about 0.047 billion gallons a day, roughly 17 billion gallons a year. That is a rounding error next to the sectors above. The larger figure activists circulate, on the order of 0.58 billion gallons a day, is mostly indirect water: which is the water used to generate the electricity the building buys. That layer is real. It is also applied in this debate exclusively to data centers, not to homes, farms, or factories that use the same grid.

How a data center is cooled decides whether it consumes water on site at all. Legacy designs cool the entire room with CRAHs, chillers, and evaporative towers to create air in the low 60s °F. Those plants evaporate continuously. New AI data centers just cool the chips.

Direct-to-chip liquid cooling (DLC) uses a closed server loop into a Cooling Distribution Unit (CDU), then a second closed facility loop to the roof. The roof plant runs dry first (fans and coils, no spray), adiabatic only on hot hours (a fraction of legacy tower water), and chillers last (energy, not makeup water). Typical useful loop temperatures are in the 80s °F; current GPU platforms are validated up to 113 °F. That is a different machine from a wet tower making 60 °F air.

The permitting fight can still be about power, land, taxes, or local process. It should not be about a national water emergency that the USGS and LBNL numbers do not show, and that the data centers now being proposed are designed to avoid.

Bringing out the knives against new data centers…literally

On July 27, 2026, anti-data-center protesters wheeled a guillotine to a public hearing in Salem, Oregon, about a proposed $5.1 billion facility.

Yes, a guillotine.

Midway through the six-hour meeting, activists confronted the developers during a break, shouting threats and insults until the company representatives left under police escort. A few days later, the city imposed a one-year moratorium on new data-center projects.

Scenes like the one in Salem are no longer rare. Across the United States, opposition to data centers has scaled rapidly. Tracking efforts by independent researchers and advocacy groups now document hundreds of local anti-data-center organizations, with some tallies exceeding 800 groups active across nearly every state and collective membership in the hundreds of thousands. Multiple public trackers and maps monitor existing facilities, proposed projects, and cancelled or rejected ones. At the same time, more than 500 local jurisdictions have enacted bans, pauses, or restrictive ordinances. Hundreds of individual projects have been delayed, withdrawn, or rejected in 2026 alone.

Activists frequently point to water consumption as a primary reason construction should be blocked in their communities. Any facility large enough to require special permitting or to attract organized opposition is almost invariably being designed for the much higher thermal densities of AI and accelerated computing. Claims often center on facilities that "use millions of gallons a day," "drain local aquifers," or "compete directly with residents and farmers for scarce water." These assertions have proven highly effective at mobilizing opposition, attracting press coverage, and then shaping local policy.

Real Numbers, Not Guesses

The figures below come from the U.S. Geological Survey's national water-use estimates (most recent comprehensive sector data) and Lawrence Berkeley National Laboratory's 2024 data-center energy and water analysis (the most detailed recent federal assessment focused on data centers).

Category	Water use (billion gallons per day)	Notes / Source
Thermoelectric power	~133	USGS 2015
Irrigation	~118	USGS 2015
Public supply	~39	USGS 2015
Industrial (self -reported)	~14.8	USGS 2015
Data centers – direct	~0.047	LBNL 2024 (~17 billion gal/year)
Data centers – indirect	~0.58	LBNL 2024 (~211 billion gal/year)

The chart below helps put the numbers above into perspective a little more clearly...

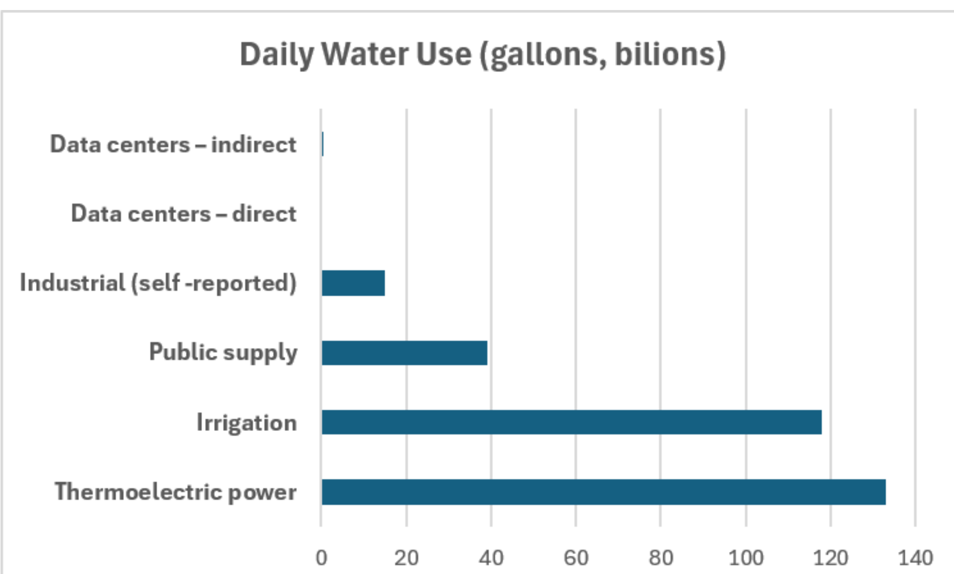

When people repeatedly hear that data centers use "nearly a billion gallons of water a day" or will "drain our water supply," the numbers sound enormous. And it does sound large—until you see it on the same scale as the rest of U.S. water use. National direct data-center water use is so small it barely registers on the chart above.

Indirect (electricity-related) water is larger, but still a thin line next to thermoelectric power, irrigation, public supply, or industry. Very few people carry the daily national water usage figures in their head; and I'm betting none of the activists promoting the scare numbers want those real figures circulated.

Direct vs. indirect water — and why the distinction matters

In the table above, **direct** means the water that flows through the data center to carry heat away from the IT equipment and out of the facility. **Indirect** means the water consumed upstream to generate the electricity the data center uses.

Data center indirect water use is roughly twelve times the direct figure. That is how 0.047 becomes about 0.58 — and how 0.58 gets talked about as nearly a billion gallons a day.

The problem is not that electricity-related water is imaginary, it's real. The problem is that this extra layer seems to be applied **only** to data centers in the public debate. The thermoelectric, irrigation, public-supply, and industrial figures shown in the table are **not** increased by the water used to generate the electricity those sectors also consume. Only the data-center total is routinely expanded this way.

If the same standard were applied evenly, every major category would grow. Add the electricity-related water to the Public Supply category, for example, and you would be counting the water used to generate power for every household, office, school, and commercial building in the U.S. The public-supply bar would jump dramatically. The same would happen to industry and, to a lesser degree, the other sectors. Comparisons would look very different.

Adding the indirect layer only to data centers alone creates an apples-to-oranges contrast that makes data center consumption appear uniquely large when, with consistent measurement standards, it is tiny.

This selective boundary is the main reason the scare numbers travel so well. Once the indirect water is stripped out, or once the same rule is applied to every sector, the claim that data centers are a dominant national water threat evaporates.

Why and how data centers use water

Data centers use water almost entirely for one purpose: removing the heat generated by servers and other IT equipment. Liquid is roughly 3,500x more efficient at holding heat vs air due to its greater density. It also transfers heat 24x faster than air. It's just much more efficient than air.

One thing that computers are great at is converting electricity to heat. The more power the equipment draws, the more heat must be removed, it's that simple.

The type of cooling system determines how much outside water—if any—is consumed in the cooling process.

Traditional, but still-common cooling designs

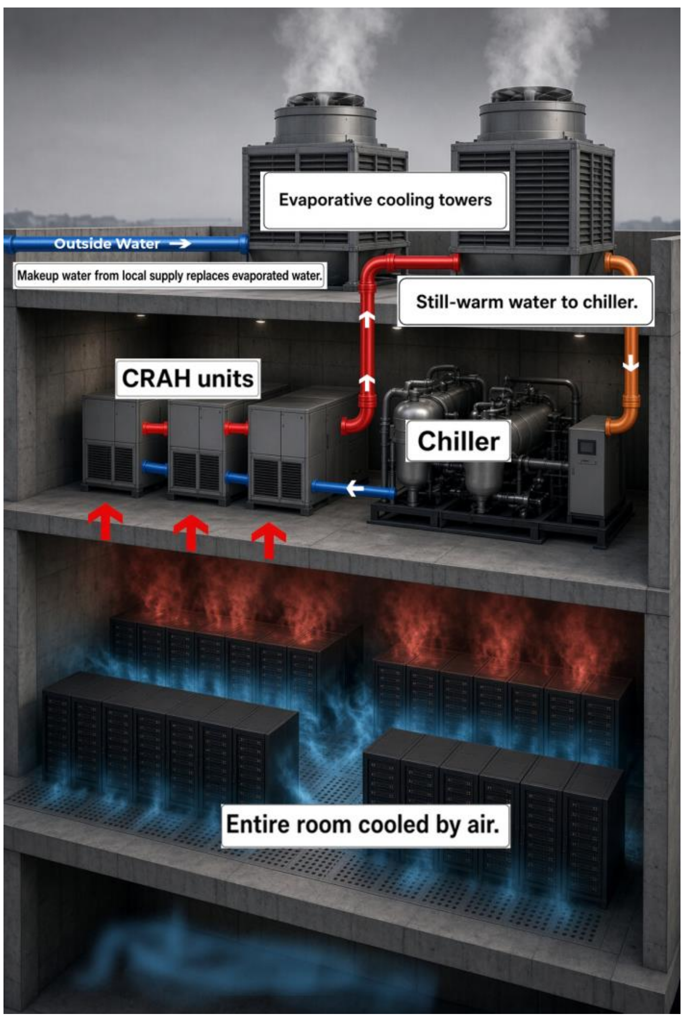

Many older and still-operating facilities rely on evaporative cooling towers, often paired with mechanical chillers that deliver low inlet temperatures, frequently in the low 60s °F.

Let's walk through the picture. The first thing to notice is that the **entire room** is being cooled rather than the components that are actually generating heat inside the servers.

Hot air rises from the servers and is drawn into the CRAH (Computer Room Air Handler) units. That air is blown across a chilled-water coil. The water in the coil absorbs the heat and is pumped toward the evaporative cooling towers on the roof.

In the towers, the warm water flows over fill media while outside air is drawn through. Some of the water evaporates; that phase change rejects heat to the atmosphere and cools the remaining water. Makeup water from the local supply replaces what was lost as vapor.

Even with the towers working, the water often isn't cold enough for the low supply-air temperatures these rooms were designed around.

That’s where the chillers come in. They work like a large refrigerator, using a refrigerant cycle to drop the water temperature further.

The fully chilled water returns to the CRAH coils. The CRAHs push air over the cold coils then deliver it into the underfloor plenum; the cool air then rises through the perforated tiles into the cold aisles. The cycle continues.

Because the system is cooling the whole room and because the towers reject heat by evaporating water, this approach consumes water continuously and uses a lot of energy, especially once IT loads climb. Modern servers, and especially GPU-dense AI systems, produce far more heat than these designs were built for.

This is what most people mean by “air cooling”: air moves through the room and the servers; water stays in the coils and on the roof. No liquid in the rack. That arrangement was built for racks of about 6–10 kW. Past roughly 20 kW, conventional room air runs out of capacity. A dense AI rack at 100 kW and up is not a hotter version of the same problem. It is a different problem, which is why the industry is moving the water to the chip.

The Modern Shift

Newer designs, especially those built for AI and accelerated computing, are moving away from continuous low-temperature evaporative cooling.

Two changes matter: operators got over their fear of liquid next to electronics, and the hardware no longer needs ice-cold air in the room.

In the late 1990s and early 2000s, direct-to-chip liquid cooling was a hobbyist trick. Hollow copper blocks bolt to the CPU and GPU. A small pump pushes treated water through the block; the warm liquid goes to a radiator and comes back cooled. That is the whole idea at chip level — and it is the same idea, at industrial scale, in a modern AI facility.

Data centers avoided it because they feared leaks. Industrial loops are built not to leak and there are fail safe mechanisms designed into the systems just in case. And the cooling task has radically changed; those rooms were designed for 6–10 kW racks. The design point now is 100 kW and up. You cannot air-cool that the old way, and you don’t need to. With DLC you cool the parts that make the heat: about 55 square inches per GPU and 12 per CPU. You are not cooling the cubic footage of the whole room.

That is why the loop can run warm. A useful inlet temperature is in the 80s °F. Current NVIDIA GPUs are validated up to 113 °F. You are not trying to hold the room in the low 60s.

Quick example: A 100,000 square foot AI facility might hold 1,650 racks, 24 servers per rack, dual CPUs and about four GPUs per server. Air cooling conditions the volume of the whole room — on the order of 2 billion cubic inches of air. DLC conditions about 10.6 million square inches of water-block surface on the chips. Same heat, far less to cool.

Closed Dual-Loop Cooling

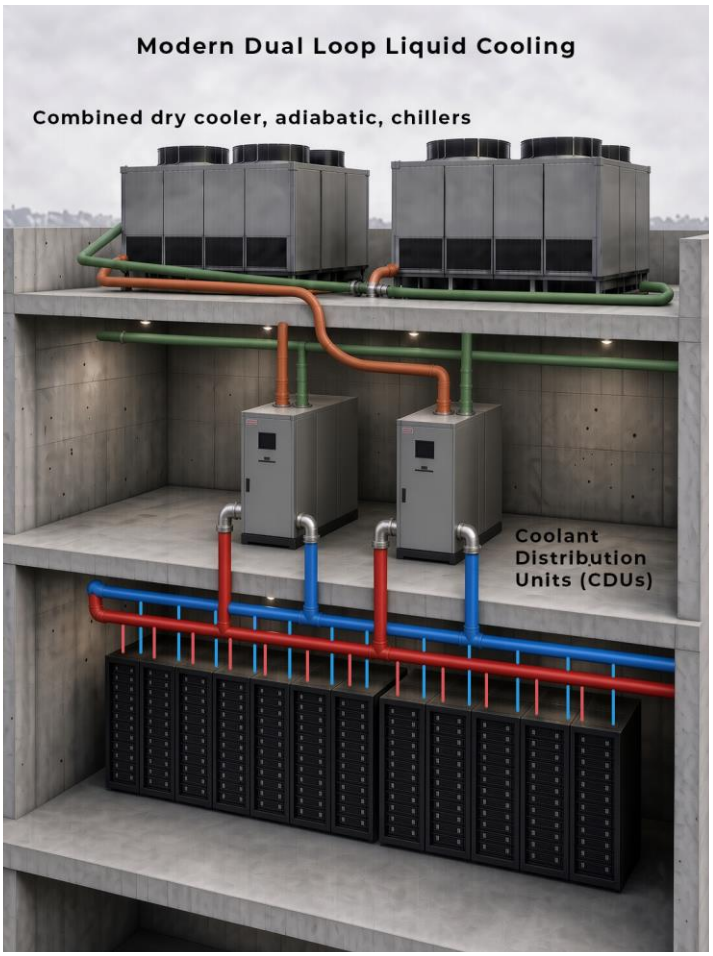

Here's how highly efficient modern data centers use water today.

There are two separate liquid loops that never mix. The server loop (with red and blue pipes/hoses) handles liquid flow for every server in the racks. The facility loop (orange and green) expels the waste heat out of the building.

On the server loop, there are liquid lines to water blocks on each CPU and GPU. The blue lines feed in cooled water and the red lines take away the heated water, delivering it to the Coolant Distribution Unit (CDU) on the next floor up.

Inside the CDU, the heat from the red lines goes through a liquid-to-liquid heat exchanger. This heat is then transferred to the facility loop which heads up to the roof.

The heat rejection units on the roof have three operating stages that work together to convert the hot water to a usable inlet temperature:

1. **Dry Cooler:** In this mode, the unit pushes air through radiators to expel heat out into the atmosphere. If your outdoor temperature is lower than the inlet temperature of your server loop, then you're using what is called "Free Cooling", meaning that the only power needed is to drive the fans expelling the heat through the radiators.

2. **Adiabatic Mode:** This is where outside water trickles down a mesh frame inside the cooler. Through evaporation, the temperature of the air pulled through the moist mesh will be significantly lower than the outside ambient temperature. The adiabatic mode uses outside water or water stored on site, but the amount of water used is just a fraction of what is used in the old-fashioned evaporative cooling we explained above.

3. **Chiller Mode:** If your ambient temperature is very high, like during a heat wave, the final stage is a refrigerant based chiller. If the first stages (dry and adiabatic) are designed correctly, chiller use should be low – only during the hottest days of the year. Chillers use more energy but do not pull any outside water.

Keep in mind that the water only needs to be cooled to around 80–85 °F to work in the server loop. That is not a heavy lift in most climates.

On-site water usage is measured as WUE — liters consumed per kilowatt-hour of IT. Open evaporative towers commonly land around 1.5 to 3 liters. A closed dual-loop plant that runs dry, and uses adiabatic only on the hot hours, sits in a different band: a few tenths of a liter, or near zero outside water if the roof can stay in dry mode. The loops do not create that result by themselves. The roof mode does.

The opposition usually describes the old plant: evaporative towers making air in the 60s. That plant will not carry a real AI load, and it is not what is being built today.

In a closed dual-loop system the water in the racks and the water in the facility pipes **stays** in those pipes. Outside water shows up only in adiabatic hours, plus a trickle for leaks and maintenance. This is a completely different consumption model than cooling towers that evaporate, then replace, water around the clock. That is the on-site water story for the AI facilities being proposed today.

Summing it up

The water argument against new data centers is built on two moves. First, take a small direct number and inflate it with electricity-related water that is not added to any other sector on the same chart. Second, describe a cooling plant that new AI facilities are leaving behind: evaporative towers and chillers manufacturing cold air for a whole room.

The new facilities proposed today are mostly DLC with two closed loops. Heat taken off a few square inches of silicon, not billions of cubic inches of an entire room. Roof units that run dry whenever the air can take the load, wet only in the heat, and mechanically chilled only when both of those are not enough. Direct water use in that design is makeup and peak adiabatic hours, not a standing tower, and thus much less than the old-style evaporation mechanism.

Indirect water still tracks the grid. If the objection is coal- or gas-fired megawatts, then protest those generation methods. If the objection is that a local water source could be drawn down over time by the facility, that is a true siting question. The plant should not sit on that source. Use non-potable water, reuse, find another supply, or another site. Work from the actual prospective water use of **that** data center — cooling mode hours, makeup rate, potable vs non-potable, etc., and back it up with data. Don't use a national scare number assigned to a building for a cooling system it won't be using.

Salem activists wheeled in a guillotine. The city then froze new projects for a year. A community can reject a data center for many valid reasons. But to reject it because the loudest voices in the room are passionately pushing false generalized numbers isn't one of them.