

TECH • A.I.

A.I. tools fueled a 34% spike in Microsoft's water consumption, and one city with its data centers is concerned about the effect on residential supply

"They were pretty secretive on what they're doing out there," says a former mayor about Microsoft's data center projects.

BY MATT O'BRIEN , HANNAH FINGERHUT , AND THE ASSOCIATED PRESS

September 09, 2023 11:01 AM EDT




Microsoft CEO Satya Nadella discusses his company's partnership with ChatGPT maker OpenAI.
Chona Kasinger/Bloomberg via Getty Images

The cost of building an artificial intelligence product like ChatGPT can be hard to measure.

Get a
welc

But one thing Microsoft-backed OpenAI needed for its technology was plenty of water, pulled from the watershed of the Raccoon and Des Moines rivers in central Iowa to cool a powerful supercomputer as it helped teach its AI systems how to mimic human writing.



As they race to capitalize on a [craze for generative AI](#), leading tech developers including Microsoft, OpenAI and [Google](#) have acknowledged that growing demand for their AI tools carries hefty costs, from expensive semiconductors to an increase in water consumption.

before a top Microsoft executive said in a speech it “was literally made next to cornfields west of Des Moines.”




PAID CONTENT

Are You Ready to Turn Your AI Strategy Into Successful Projects?



FROM HEWLETT PACKARD ENTERPRISE

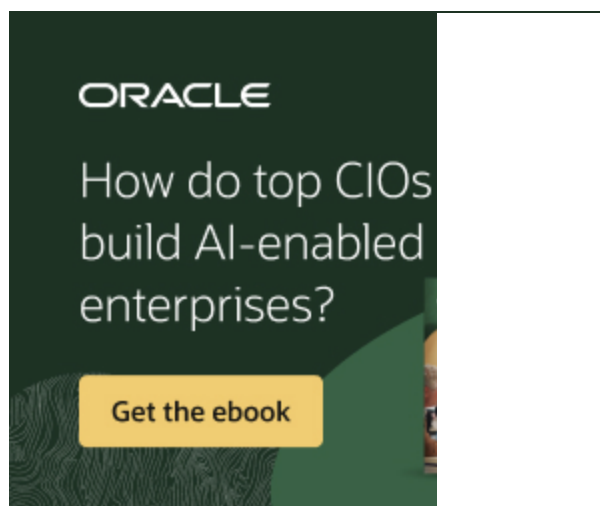
Building a large language model requires analyzing patterns across a huge trove of human-written text. All of that computing takes a lot of electricity and generates a lot of heat. To keep it cool on hot days, data centers need to pump in water — often to a cooling tower outside its warehouse-sized buildings.



In its [latest environmental report](#), Microsoft disclosed that its global water consumption spiked 34% from 2021 to 2022 (to nearly 1.7 billion gallons, or more than 2,500 Olympic-sized swimming pools), a sharp increase compared to previous

“It’s fair to say the majority of the growth is due to AI,” including “its heavy investment in generative AI and partnership with OpenAI,” said Shaolei Ren, a researcher at the University of California, Riverside who has been trying to calculate the environmental impact of generative AI products such as ChatGPT.


In a paper due to be published later this year, Ren’s team estimates ChatGPT gulps up 500 milliliters of water (close to what’s in a 16-ounce water bottle) every time you ask it a series of between 5 to 50 prompts or questions. The range varies depending on where its servers are located and the season. The estimate includes indirect water usage that the companies don’t measure — such as to cool power plants that supply the data centers with electricity.



“Most people are not aware of the resource usage underlying ChatGPT,” Ren said. “If you’re not aware of the resource usage, then there’s no way that we can help conserve the resources.”

Google [reported](#) a 20% growth in water use in the same period, which Ren also largely attributes to its AI work. Google’s spike wasn’t uniform — it was steady in Oregon where its [water use has attracted](#) public attention, while doubling outside Las Vegas. It was also thirsty in Iowa, drawing more potable water to its Council Bluffs data centers than anywhere else.

In response to questions from The Associated Press, Microsoft said in a statement this week that it is investing in research to measure AI’s energy and carbon footprint “while working on ways to make large systems more efficient, in both training and application.”



“We will continue to monitor our emissions, accelerate progress while increasing our use of clean energy to power data centers, purchasing renewable energy, and other efforts to meet our sustainability goals of being carbon negative, water positive and zero waste by 2030,” the company’s statement said.

OpenAI echoed those comments in its own statement Friday, saying it’s giving “considerable thought” to the best use of computing power.


“We recognize training large models can be energy and water-intensive” and work to improve efficiencies, it said.



Microsoft made its first \$1 billion investment in San Francisco-based OpenAI in 2019, more than two years before the startup introduced ChatGPT and sparked worldwide [fascination with AI advancements](#). As part of the deal, the software giant would supply computing power needed to train the AI models.

power its cloud computing services for more than a decade. Its fourth and fifth data centers are due to open there later this year.

“They’re building them as fast as they can,” said Steve Gaer, who was the city’s mayor when Microsoft came to town. Gaer said the company was attracted to the city’s commitment to building public infrastructure and contributed a “staggering” sum of money through tax payments that support that investment.



“But, you know, they were pretty secretive on what they’re doing out there,” he added.

Microsoft first said it was developing one of the world’s most powerful supercomputers for OpenAI in 2020, declining to reveal its location to AP at the time but describing it as a “single system” with more than 285,000 cores of conventional semiconductors, and 10,000 graphics processors — a kind of [chip that’s become crucial](#) to AI workloads.

Experts have said it can make sense to “pretrain” an AI model at a single location because of the large amounts of data that need to be transferred between computing cores.

It wasn't until late May that Microsoft's president, Brad Smith, disclosed that it had built its "advanced AI supercomputing data center" in Iowa, exclusively to enable OpenAI to train what has become its fourth-generation model, GPT-4. The model now powers premium versions of ChatGPT and some of Microsoft's own products and has accelerated a debate about containing AI's societal risks.

"It was made by these extraordinary engineers in California, but it was really made in Iowa," Smith said.

In some ways, West Des Moines is a relatively efficient place to train a powerful AI system, especially compared to Microsoft's data centers in Arizona that consume far more water for the same computing demand.

"So if you are developing AI models within Microsoft, then you should schedule your training in Iowa instead of in Arizona," Ren said. "In terms of training, there's

For much of the year, Iowa's weather is cool enough for Microsoft to use outside air to keep the supercomputer running properly and vent heat out of the building. Only when the temperature exceeds 29.3 degrees Celsius (about 85 degrees Fahrenheit) does it withdraw water, the company has said in a public disclosure.

That can still be a lot of water, especially in the summer. In July 2022, the month before [OpenAI says it completed](#) its training of GPT-4, Microsoft pumped in about 11.5 million gallons of water to its cluster of Iowa data centers, according to the West Des Moines Water Works. That amounted to about 6% of all the water used in the district, which also supplies drinking water to the city's residents.

In 2022, a document from the West Des Moines Water Works said it and the city government "will only consider future data center projects" from Microsoft if those projects can "demonstrate and implement technology to significantly reduce peak water usage from the current levels" to preserve the water supply for residential and other commercial needs.

Microsoft said Thursday it is working directly with the water works to address its feedback. In a written statement, the water works said the company has been a good partner and has been working with local officials to reduce its water footprint while still meeting its needs.

--

O'Brien reported from Providence, Rhode Island.