Informational Materials - Speaker Handouts

- Meeting Summary: Mapping out how machine learning and artificial intelligence will change Great Lakes observations, modeling, and forecasting in the coming decade
- Executive Summary: 2024 United States Data Center Energy Usage Report
- Fact sheet: Aqueous film-forming foam
- Fact sheet: Building a circular water economy in the Great Lakes
- Fact sheet: Cleveland Water Alliance
- Fact sheet: Water reuse solutions for data centers

BAMS
Meeting Summary

Mapping Out How Machine Learning and Artificial Intelligence Will Change Great Lakes Observations, Modeling, and Forecasting in the Coming Decade

Dani Jones,[a] Scott Steinschneider,[b] Paul Roebber,[c] Sage Osborne,[d] Lauren Fry,[e] Lacey Mason,[e] Andrea Vander Woude,[e] Mantha S. Phanikumar,[f] Nathan Fox,[g] William S. Currie,[h] Silvia Santa Maria Newell,[h,i] Jia Wang,[e] Alisa Young,[e] Lindsay Fitzpatrick,[a] Yi Hong,[a] Hazem Abdelhady,[a] William J. Pringle,[j] E. Anders Kiledal,[k] and Andrew D. Gronewold[h]

KEYWORDS:
Inland seas/lakes;
Watersheds;
Forecasting;
Water resources;
Artificial intelligence;
Machine learning

AI Horizons

What: Twenty-two scientists gathered to explore how machine learning and artificial intelligence could transform Great Lakes observations, modeling, and forecasting. Participants discussed enhancing predictive models, improving observational network design, and identifying key challenges and opportunities. They aimed to establish a framework for a Great Lakes ML/AI community of practice and to develop a strategic vision for the next decade.

When: 22–23 July 2024

Where: Ann Arbor, Michigan

DOI: 10.1175/BAMS-D-24-0304.1

Corresponding author: Dani Jones, dannes@umich.edu

In final form 2 December 2024

AFFILIATIONS: [a] Cooperative Institute for Great Lakes Research, University of Michigan, Ann Arbor, Michigan; [b] Department of Biological and Environmental Engineering, Cornell University, Ithaca, New York; [c] Atmospheric Science Group, Department of Mathematical Sciences, University of Wisconsin–Milwaukee, Milwaukee, Wisconsin; [d] NOAA/Pacific Marine Environmental Laboratory, Seattle, Washington; [e] NOAA/Great Lakes Environmental Research Laboratory, Ann Arbor, Michigan; [f] Department of Civil and Environmental Engineering, Michigan State University, East Lansing, Michigan; [g] Michigan Institute for Data and AI in Society, University of Michigan, Ann Arbor, Michigan; [h] School for Environment and Sustainability, University of Michigan, Ann Arbor, Michigan; [i] Michigan Sea Grant, Ann Arbor, Michigan; [j] Argonne National Laboratory, Lemont, Illinois; [k] Department of Earth and Environmental Sciences, University of Michigan, Ann Arbor, Michigan

1. Introduction

The Michigan Institute for Data and AI in Society (MIDAS) and the Cooperative Institute for Great Lakes Research (CIGLR) hosted a summit in Ann Arbor, Michigan, from 22 to 23 July 2024. The event brought together 22 scientists from across the country to envision how machine learning and artificial intelligence might be used to address some of the most pressing challenges facing Great Lakes science, restoration, and management over the next decade.

2. Background

The Laurentian Great Lakes—comprising Lakes Superior, Michigan, Huron, Erie, and Ontario—are the largest group of freshwater lakes in the world by the total area, holding approximately 21% of Earth's surface freshwater. In addition to their considerable size, the Great Lakes serve as a critical resource for over 40 million people in the United States and Canada. They provide drinking water, support commercial and recreational fishing, facilitate transportation and commerce, and offer opportunities for tourism, recreation, and connecting with nature. The Great Lakes also enrich regional biodiversity by supporting a wide range of interconnected ecosystems. Managed cooperatively by two neighboring countries, this dynamic system represents a case study for international water and ecosystem management and serves as a key testbed for understanding aquatic, terrestrial, and climatological systems globally.

Many different local, state, federal, and tribal government agencies, nonprofit organizations, academic institutions, private entities, and community stakeholders are involved in the complex scientific and management landscape of the Great Lakes. For example, the International Joint Commission (IJC) coordinates water management and water quality efforts between the United States and Canada, reflecting their mutual responsibilities for water use and safety. A variety of federal agencies, regional consortia, and national laboratories have developed regional, state-of-the-art monitoring networks that include buoys, vessel fleets, flux towers, and remote sensing capabilities. Notable organizations include the U.S. Environmental Protection Agency, Environment and Climate Change Canada, the National Oceanic and Atmospheric Administration's (NOAA) Great Lakes Environmental Research Laboratory (GLERL), the Cooperative Institute for Great Lakes Research (CIGLR), and the Great Lakes Observing System (GLOS). Initiatives such as the Great Lakes Restoration Initiative, CoastWatch, the Synthesis, Observations, and Response System (SOAR), and Submerged Aquatic Vegetation (SAV) Mapping also play significant roles. The high concentration of observational networks in the region sets some parts of the Great Lakes apart as some of the best-monitored, large open water regions in the world. These efforts underscore the collaborative approach taken by multiple stakeholders to address pressing challenges facing the Great Lakes.

For example, despite considerable progress in recent years, accurate forecasting for the Great Lakes region remains a key challenge. Changes in lakewide water levels are set by a small residual of several large fluxes that can be difficult to forecast accurately (i.e., evaporation, precipitation, runoff). Even small errors in the estimates of these fluxes can lead to considerable errors in the water level forecast. As climate change continues to alter atmospheric thermodynamics and dynamics, projecting the net effect on the Great Lakes remains challenging. Evaporation tends to lower water levels and precipitation tends to raise them—both processes will be affected by climate change, and it is far from clear which one will dominate within a given forecast period. Accurate forecasting of ice conditions and harmful algal blooms (HABs) are also challenging.

Many Great Lakes datasets and studies are watershed specific and small scale, but understanding the entire Great Lakes—much like any large basin on Earth—requires a more holistic, integrated approach that considers the Great Lakes as a single system. Machine learning (ML) and artificial intelligence (AI), underpinned by the necessary data and infrastructure, present a unique opportunity for considerable advancements in forecasting, as well as other areas such as process-based understanding and observing network design.

Machine learning and artificial intelligence are transforming environmental science. For example, the field of oceanography has been revolutionized by ML/AI, both in terms of modeling capacity and scientific understanding (Sonnewald et al. 2021). Atmospheric science is similarly evolving, especially in the weather forecasting industry (Eyring et al. 2024). For example, the European Centre for Medium-Range Weather Forecasts (ECMWF) now runs a data-driven forecast system (Lang et al. 2024). ML/AI could also help improve relevant large-scale theories for hydrologic simulations and forecasting (Nearing et al. 2021). The Great Lakes research and management community should engage with these emerging toolkits in a coordinated and systematic manner to effectively achieve the greatest impact. Given the complexity of the science involved and the multifaceted nature of the Great Lakes community, we need a strategic approach.

3. CIGLR summit: AI horizons

The "AI Horizons" summit represents a strategic effort to envision the integration of ML/AI into Great Lakes science, management, and restoration. CIGLR Summits convene groups of 20–30 invited experts meeting for 2–3 days to summarize the state of knowledge and recommend future directions on Great Lakes problems that span decadal time scales. On 22 and 23 July 2024, CIGLR hosted a summit on the University of Michigan campus, specifically in the Samuel T. Dana Building, which is home to the School for Environment and Sustainability (SEAS). The summit brought together 22 researchers from across the country. The organizers identified two primary goals for the summit: creating a framework for a Great Lakes ML/AI community of practice and developing a decadal vision on how ML/AI could transform Great Lakes science. Over the 2-day summit, participants made progress on scaffolding and drafting up the vision document and engaged in foundational conversations on building a Great Lakes–centric community of practice.

The first day of the summit began with scene-setting talks. The organizers emphasized the importance of collaboration, stating that the success of the summit, like any collaborative endeavor, relies on the collective experience, energy, focus, and goodwill of the participants. To harness the collective intelligence of the participants, organizers implemented several tools and practices:

- Shared, online notes documents that all summit participants could edit.
- A group Zotero library, allowing for group management of relevant literature, citations, and bibliographies.
- Dynamically defined working groups that evolved throughout the summit in response to participant feedback.
- A “parking lot” for noting important topics that fell outside the formal agenda but warranted future discussion.
- Web-based Q&A polls throughout the summit via the “Slido” tool.
- Dedicated blocks of writing time, as part of the program, to allow for informal collaborative conversation, individual focus, and codevelopment of text to articulate ideas.

Jing Liu, the Executive Director of MIDAS, gave a brief overview of AI in science and engineering. She noted that data and AI are becoming essential tools in many domains, transforming both disciplinary and interdisciplinary research. This trend calls for new support mechanisms, such as new institutes. As one such institute, MIDAS is currently focused on enabling the adoption of AI methods in various research domains, promoting responsible and ethical AI, and exploring new ways of cross-sector collaboration with academia, industry, and government. Liu highlighted that the current academic research model has struggled to implement AI at the needed scale and pace to address pressing scientific problems, particularly in comparison to industry, where resources and talent tend to be concentrated.

In a second scene-setting talk, Dani Jones (CIGLR) echoed definitions of ML and AI as formulated by Liu. Specifically, “AI” refers to machines and/or algorithms that can perform tasks that normally require human intelligence, such as learning, reasoning, using and comprehending language, decision-making, and so on. Under that definition, “ML” may be considered a subset of AI focused on data-driven learning and prediction, which does have some overlap with traditional statistics.

Jones gave a brief overview of how machine learning has shaped the field of Earth system science. In oceanography, there have been considerable advances in subgrid scale parameterization and equation discovery (Bolton and Zanna 2019; Zanna and Bolton 2020), novel data analyses using unsupervised learning (Sonnewald et al. 2019, 2020), and hypothesis generation (Sonnewald et al. 2021). To set the stage to consider a framework that would lead into a Great Lakes–centric AI/ML community of practice, Jones shared examples of other centers of practice. In atmospheric science, organizations such as the Cooperative Institute for Research in the Atmosphere (CIRA) have pioneered the use of ML/AI in detection, inference, and estimation from satellite-based remote sensing products (Lagerquist and Ebert-Uphoff 2022). Large-scale initiatives and institutes like the NSF AI Institute for Research on Trustworthy AI in Weather, Climate, and Coastal Oceanography (McGovern et al. 2024) and the NASA JPL Science Understanding through Data Science (SUDS) initiative work to integrate ML/AI into Earth system sciences, emphasizing domain knowledge, trustworthiness, interpretability, and explainability.

Jones emphasized that Great Lakes science and management is just starting to use ML and is still in the early stages of conceptualizing how AI might be used. There have been ML-driven advances in predicting wave heights (Hu et al. 2021), ice cover (Abdelhady and Troy 2024), and water levels. There have also been advances in autonomous underwater sampling (Zhang et al. 2024) and model emulation and uncertainty quantification (Pringle et al. 2024). At present, there are efforts at GLERL and CIGLR to use data-driven methods to improve water level forecasting and improve understanding of the drivers of HABs. There are undoubtedly more efforts in the research community of which the authors are unaware, emphasizing the need for a central hub to coordinate these efforts and integrate ML/AI into Great Lakes science more effectively, supporting research, the exploration of new methods, training, and other relevant knowledge sets.

Dynamic working groups. Before the summit, the organizers proposed the following working groups: 1) modeling and forecasting capacity, 2) observing network design, 3) data pipelines and data assimilation, and 4) hazards and risks. An initial survey revealed that interest in these working groups was unbalanced. The overwhelming majority of participants selected "modeling & forecasting capacity" as their first choice, while there was very little interest in "hazards and risks." In response, the organizers split the modeling & forecasting capacity group into two: one focusing on predictive modeling and the other on operationalization and applications. The initially proposed "observing network design" and "data pipelines and data assimilation" groups were combined into a single group, and a new "generative AI" group was established. Participants then self-sorted into the working groups. During the morning working group discussions, participants focused on their respective topics. The groups highlighted issues, identified opportunities, and made recommendations to improve Great Lakes research and management through ML/AI:

- Modeling and Forecast Capacity A: Predictive Models and Techniques. This group focused on pathways to develop and refine predictive models using machine learning, address technical challenges, enhance model accuracy, and explore innovative methodologies. Key discussion points included the following:
  - Consideration of data from a statistical point of view: independence of samples, sufficiently large spatial datasets, and time series.
  - Opportunities for linking highly time-resolved mechanistic models with less frequently sampled biological data to improve beach closures and understanding of harmful algal blooms. The group also emphasized leveraging physics-informed models while addressing the issue of "black box" models and various data characteristics depending on the application, such as dimensionality, sampling frequency, and predictive time scales.
  - Challenges such as overfitting, enhancing model accuracy, and trade-offs between accuracy and generalizability.

- Modeling and Forecast Capacity B: Operationalization and Applications. This group addressed the practical implementation of ML approaches in real-world scenarios, supporting decision-making processes and engaging stakeholders. Discussion points included the following:
  - Forecasting at various time scales (short-term, subseasonal to seasonal, and years to decades).
  - Using ML to better reconstruct past and current conditions (hindcast/nowcasting).
  - Supporting large-scale infrastructure operation and policy with reinforcement learning and global approximators.
  - Key challenges included stakeholder engagement and building trust while explaining and transparently communicating ML-based operations to the public and operators.

- Generative AI: This group focused on exploring the potential of generative AI in Great Lakes science, restoration, and management. They discussed the following:
  - The strengths and weaknesses of current technologies, such as integrating physics into cost functions.

    - Potential for generative AI to fill in data gaps, considering that AI might not extrapolate well due to reliance on training datasets.
    - Highlighting how physics-based optimization could help capture extremes and incorporate domain expertise.

- Data Pipelines and Observing Network Design: This group looked at enhancing data integration and access, data ingestion pipelines, and mechanisms for AI to answer relevant questions for observation systems. They emphasized the following:
    - The importance of accessible, robust, and analysis-ready data for ML/AI development.
    - Utilizing AI for data curation, quality control, creating metadata, and supporting public decision-making through natural language processing.
    - Addressing data sparsity, temporal issues, and ensuring accurate AI constraints by leveraging physics-informed models.

Following the morning working group breakout discussions, the participants reconvened to report on their discussions. It was decided that the working groups should be reorganized dynamically again by the participants before the afternoon session. The reorganized working groups aligned with the planned sections of the paper, specifically the following:

- Philosophy, context, and problem formulation: Refining research questions, theoretical frameworks, and contextualizing challenges unique to the Great Lakes.
- Benefits of AI in addressing Great Lakes challenges: Exploring specific applications and improvements brought by ML/AI.
- Connecting with the wider landscape, best practices, and interoperability: Ensuring alignment with global initiatives and developing best practices.

The second day of the summit was dedicated to drafting the perspectives manuscript and developing the community of practice framework. Short breakout sessions allowed teams to refine sections of the manuscript, integrating insights from the previous day. Alternating between full-group discussions and focused writing sessions helped ensure alignment and coherence across different sections.

By late morning, outlines of key sections were ready for review. The generative AI group contributed insights into the strengths and weaknesses of current ML/AI technologies in meteorology, emphasizing the rapid growth in this discipline. They noted the potential for tailoring cost functions to capture physics better and using AI to extend data in data-sparse regions by leveraging physical laws. This approach motivates the collection of more data to enhance AI capabilities. Current technologies often struggle with capturing extremes due to issues with extrapolation, which is crucial for climate change projections. Physics-based optimizations can help incorporate domain expertise and potentially address these challenges.

A collaborative editing session followed, where participants provided feedback and integrated complementary ideas. The summit concluded with final remarks, reinforcing the importance of continued collaboration and outlining the next steps for the Great Lakes AI Laboratory and planned publications.

4. Summit outcomes

The summit aims to ultimately produce a perspectives paper on the effective integration of ML/AI into Great Lakes science and management. The work on the perspectives paper, which started during the summit, has continued post-summit, with participants engaging in writing groups and using collaborative tools to ensure coherence and comprehensive coverage of the identified themes and references.

a. Building up the Great Lakes AI community. The second goal of the summit was to develop strategies, tools, and practices for fostering a collaborative AI research community in the Great Lakes region. As highlighted during the summit discussions, this effort should include the following:

- A framework for a Great Lakes ML/AI community of practice: Establishing a structured approach to foster accessible, analysis-ready data [e.g., drawing inspiration from community initiatives like Pangeo (Odaka et al. 2020)] and reproducibility of both code and datasets. It was mentioned that the Great Lakes research community may not have the scale needed to create its own large-scale infrastructure; it may be more beneficial to join existing larger efforts, introducing a specific subfocus on the Great Lakes.
- Community engagement: Promoting participation and collaboration across different stakeholders, leveraging diverse expertise and perspectives. This could take the form of sessions at national and international conferences (e.g., the American Geophysical Union's Fall Meeting).
- Balancing act: Addressing community-building efforts while mitigating concerns of communication burnout. This includes respecting the limited bandwidth of individuals and prioritizing sustainable development.
- Support structures: Implementing sustainable support structures to maintain engagement and ensure continuous collaboration. This includes developing mechanisms for ongoing communication, resource sharing, and collaborative efforts.

b. Stakeholder engagement and trust building. For successful integration, there must be a focus on engaging stakeholders and building trust in ML/AI systems. Transparent communication and explaining model outputs effectively are critical for gaining public and stakeholder trust.

5. Next steps

a. Perspectives manuscript. Summit participants will draft and submit a perspectives article on the future of ML and AI in Great Lakes science and management. This manuscript, being developed by all summit attendees with additional expert input, will serve as a foundational document.

b. Launch of the Great Lakes AI laboratory. The newly launched Great Lakes AI Laboratory will serve as a focal point for integrating ML and AI into Great Lakes research, management, and restoration, addressing the pressing need for advanced data pipelines, analytics, and predictive modeling in the region. Although CIGLR—consisting of a research institute and a Regional Consortium of academic, nongovernmental organization, and private sector partners—and NOAA GLERL will serve as the initial hosts for this initiative, the Great Lakes AI Laboratory is designed to be an open, collaborative community. We invite anyone with research interests or a stake in Great Lakes science, management, and restoration, as well as environmental machine learning experts looking to engage with the Great Lakes community, to join us. Specific platforms and events that the Great Lakes AI Laboratory will use may include a website, Slack channel, GitHub organization, mailing list, open workshops, and hackathons. By developing advanced predictive models, facilitating data integration, providing training and resources, and encouraging interdisciplinary collaboration, the Great Lakes AI Laboratory seeks to bridge the gap between data science and environmental science in the Great Lakes, fostering innovation and enhancing the region's capacity to address its unique challenges. We invite interested parties to join us as part of the Great Lakes AI Laboratory GitHub organization (https://github.com/great-lakes-ai-lab).

Acknowledgments. Funding for the summit was provided by the Cooperative Institute for Great Lakes Research (CIGLR), which is funded through the NOAA Cooperative Agreement with the University of Michigan (NA22OAR4320150). ADG (Gronewold) was partially supported through the NSF Global Centers program (NSF Award 2330317). This is NOAA GLERL Contribution Number 2067 and CIGLR Contribution Number 1254. The authors wish to thank all summit participants for their time, focus, and perspectives. We extend a special thanks to Margaret Throckmorton and Mary Ogdahl (CIGLR) and Beth Ubserseder (MIDAS) for their invaluable support in organizing and carrying out the meeting.

References

Abdelhady, H., and C. Troy, 2024: A deep learning approach for modeling and hindcasting Lake Michigan ice cover. arXiv, 2407.04937v1, https://doi.org/10.48550/arXiv.2407.04937.

Bolton, T., and L. Zanna, 2019: Applications of deep learning to ocean data inference and subgrid parameterization. *J. Adv. Model. Earth Syst.*, **11**, 376–399, https://doi.org/10.1029/2018MS001472.

Eyring, V., and Coauthors, 2024: Pushing the frontiers in climate modelling and analysis with machine learning. *Nat. Climate Change*, **14**, 916–928, https://doi.org/10.1038/s41558-024-02095-y.

Hu, H., A. J. van der Westhuysen, P. Chu, and A. Fujisaki-Manome, 2021: Predicting Lake Erie wave heights and periods using XGBoost and LSTM. *Ocean Modell.*, **164**, 101832, https://doi.org/10.1016/j.ocemod.2021.101832.

Lagerquist, R., and I. Ebert-Uphoff, 2022: Can we integrate spatial verification methods into neural network loss functions for atmospheric science? *Artif. Intell. Earth Syst.*, **1**, e220021, https://doi.org/10.1175/AIES-D-22-0021.1.

Lang, S., and Coauthors, 2024: AIFS – ECMWF's data-driven forecasting system. arXiv, 2406.01465v2, https://doi.org/10.48550/arXiv.2406.01465.

McGovern, A., and Coauthors, 2024: AI2ES: The NSF AI Institute for research on trustworthy AI for weather, climate, and coastal oceanography. *AI Mag.*, **45**, 105–110, https://doi.org/10.1002/aaai.12160.

Nearing, G. S., F. Kratzert, A. K. Sampson, C. S. Pelissier, D. Klotz, J. M. Frame, C. Prieto, and H. V. Gupta, 2021: What role does hydrological science play in the age of machine learning? *Water Resour. Res.*, **57**, e2020WR028091, https://doi.org/10.1029/2020WR028091.

Odaka, T. E., A. Banihirwe, G. Eynard-Bontemps, A. Ponte, G. Maze, K. Paul, J. Baker, and R. Abernathey, 2020: The Pangeo Ecosystem: Interactive computing tools for the geosciences: Benchmarking on HPC. *Tools and Techniques for High Performance Computing*, G. Juckeland and S. Chandrasekaran, Eds., Springer International Publishing, 190–204.

Pringle, W. J., and Coauthors, 2024: Coupled lake-atmosphere-land physics uncertainties in a Great Lakes regional climate mode. ESS Open Archive, 171052472.24971474/v1, https://doi.org/10.22541/essoar.171052472.24971474/v1.

Sonnewald, M., C. Wunsch, and P. Heimbach, 2019: Unsupervised learning reveals geography of global ocean dynamical regions. *Earth Space Sci.*, **6**, 784–794, https://doi.org/10.1029/2018EA000519.

——, S. Dutkiewicz, C. Hill, and G. Forget, 2020: Elucidating ecological complexity: Unsupervised learning determines global marine eco-provinces. *Sci. Adv.*, **6**, eaay4740, https://doi.org/10.1126/sciadv.aay4740.

——, R. Lguensat, D. C. Jones, P. D. Dueben, J. Brajard, and V. Balaji, 2021: Bridging observations, theory and numerical simulation of the ocean using machine learning. *Environ. Res. Lett.*, **16**, 073008, https://doi.org/10.1088/1748-9326/ac0eb0.

Zanna, L., and T. Bolton, 2020: Data-driven equation discovery of ocean mesoscale closures. *Geophys. Res. Lett.*, **47**, e2020GL088376, https://doi.org/10.1029/2020GL088376.

Zhang, Y., and Coauthors, 2024: Using a long-range autonomous underwater vehicle to find and sample harmful algal blooms in Lake Erie. *Limnol. Oceanogr. Methods*, **22**, 473–483, https://doi.org/10.1002/lom3.10621.

LBNL-2001637

BERKELEY LAB

Energy Analysis & Environmental Impacts Division

2024 United States Data Center Energy Usage Report

Arman Shehabi, Sarah J. Smith, Alex Hubbard, Alex Newkirk, Nuoa Lei, Md Abu Bakar Siddik, Billie Holecek, Jonathan Koomey, Eric Masanet, and Dale Sartor
Energy Analysis and Environmental Impacts Division, Lawrence Berkeley National Laboratory

December 2024

This work was supported by the DOE Office of Energy Efficiency and Renewable Energy (EERE), Office of Industrial Efficiency and Decarbonization Office (IEDO) under Lawrence Berkeley National Laboratory Contract No. DE-AC02-05CH11231.

Disclaimer

Copyright Notice

Citation

Shehabi, A., Smith, S.J., Hubbard, A., Newkirk, A., Lei, N., Siddik, M.A.B., Holecek, B., Koomey, J., Masanet, E., Sartor, D. 2024. 2024 United States Data Center Energy Usage Report. Lawrence Berkeley National Laboratory, Berkeley, California. LBNL-2001637

Acknowledgements

We would also like to acknowledge and thank the expert reviewers from approximately 50 organizations representing major data center equipment manufacturers and data center end users/owners that took time to answer questions from the authors and provided comments on draft versions of the report.

Input data for this report was provided by Omdia Research, the Dell'Oro Group, S&P Global, and the International Data Corporation.

The research reported in this report was conducted by Lawrence Berkeley National Laboratory with support from the Department of Industrial Efficiency and Decarbonization Office. Lawrence Berkeley National Laboratory is supported by the Office of Science of the United States Department of Energy and operated under Contract Grant No. DE-AC02-05CH11231.

Table of Contents

List of Figures

List of Tables

Executive Summary

The Energy Act of 2020 calls for the U.S. Department of Energy to make available to the public an update to Lawrence Berkeley National Laboratory's prior study entitled *United States Data Center Energy Usage Report* (2016). This report, designed to meet that Congressional request, estimates historical data center electricity consumption back to 2014, relying on previous studies and historical shipment data. This report also provides a scenario range of future demand out to 2028 based on new trends and the most recent available data. Figure ES-1 (below) provides an estimate of total U.S. data center electricity use including servers, storage, network equipment, and infrastructure from 2014 through 2028.

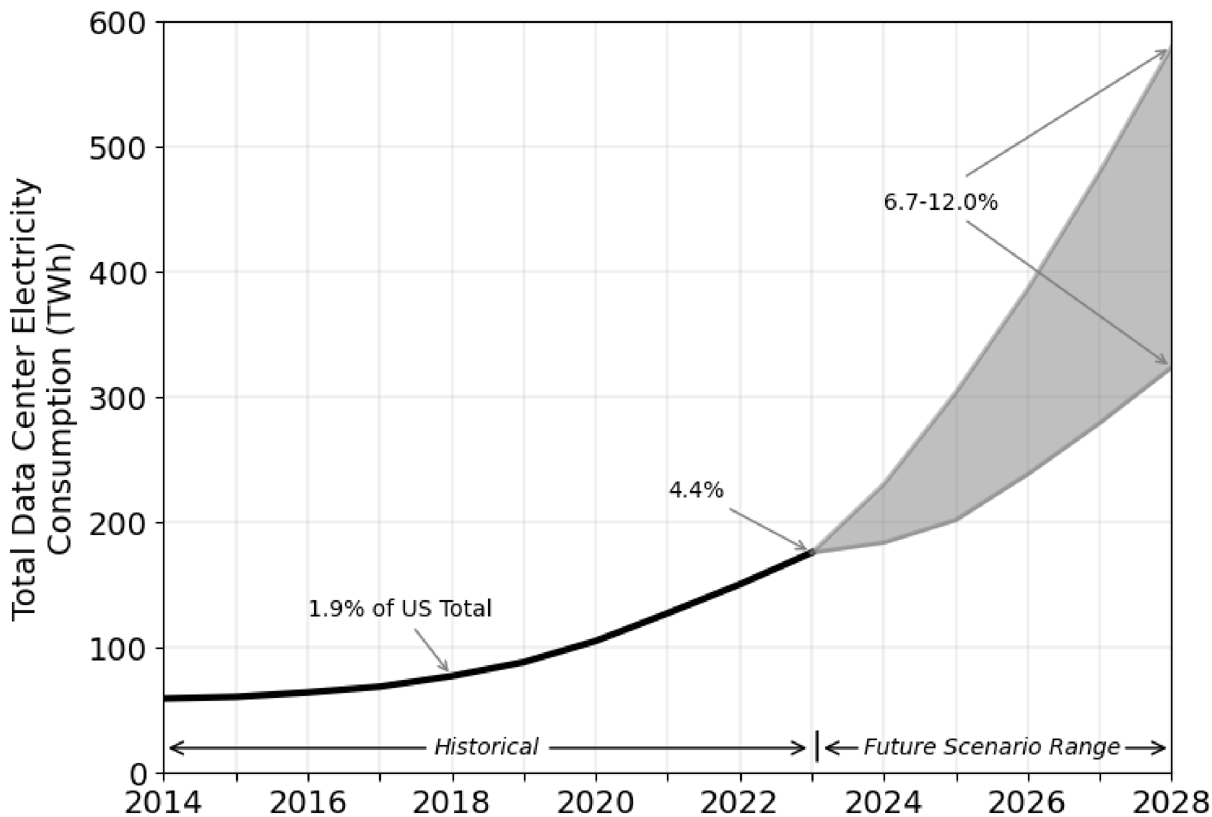

Figure ES-1. Total U.S. data center electricity use from 2014 through 2028.

As Figure ES-1 shows, U.S. data center annual energy use remained stable between 2014–2016 at about 60 TWh, continuing a minimal growth trend observed since about 2010. In 2017, the overall server installed base started growing and Graphic Processing Unit (GPU)-accelerated servers for artificial intelligence (AI) became a significant enough portion of the data center server stock that total data center electricity use began to increase again, such that by 2018 data centers consumed about 76 TWh, representing 1.9% of total annual U.S. electricity consumption. U.S. data center energy use has continued to grow at an increasing rate, reaching 176 TWh by 2023, representing 4.4% of total U.S. electricity consumption.

With significant changes observed in the data center sector in recent years, owing to the rapid emergence of AI hardware, total data center energy use after 2023 is presented as a range to reflect various scenarios. These scenarios capture ranges of future equipment shipments and operational practices, as well as variations in cooling energy use. The equipment variations are based on the assumed number of GPUs shipped each year, which depends on the future GPU demand and the ability of manufacturers to meet those demands. Average operational practices for GPU-accelerated servers represent how much computational power, and how often AI hardware in the installed base is used, to meet AI workload demand. Cooling energy use variations are based on scenarios in cooling system selection type and efficiency of those cooling systems, such as shifting to liquid base cooling or moving away from evaporative cooling. Together, the scenario variations provide a range of total data center energy estimates, with the low and high end of roughly 325 and 580 TWh in 2028, as shown in Figure ES-1. Assuming an average capacity utilization rate of 50%, this annual energy use range would translate to a total power demand for data centers between 74 and 132 GW. This annual energy use also represents 6.7% to 12.0% of total U.S. electricity consumption forecasted for 2028.

Historically, data center electricity use increased substantially from 2000–2005, roughly doubling during that period. During the early and mid-2010s, a shift from on-premise data centers to colocation or cloud facilities helped enable efficiency improvements that allowed data center electricity use to remain nearly constant at a time when the data center industry grew significantly, with a large expansion of data center services. The efficiency strategies that allowed the industry to avoid increased energy needs during this period included improved cooling and power management, increased server utilization rates, increased computational efficiencies, and reduced server idle power.

While many of these efficiency strategies continue to provide significant energy efficiency improvements in data center design and operation, the expansion of data center services into areas that require new types of hardware has ended the era of generally flat data center energy use. Most notably, the rapid growth in accelerated servers has caused current total data center energy demand to more than double between 2017 and 2023, and continued growth in the use of accelerated servers for AI services could cause further substantial increases by the end of this decade. The current and possible near-future surge in energy demand highlights the need for future research to understand the early-stage, rapidly changing AI segment of the data center industry and identify new efficiency strategies to minimize the resource impacts of this growing and increasingly significant sector in our economy.

Areas of future research identified in this report include benchmarking initiatives, collaborations with electric utilities, and technology development, all of which would be furthered by greater transparency in data center energy use, as the lack of data availability significantly limits the analysis in this report. The estimates in this report are based on a “bottom-up” energy use model that calculates total electricity use from an installed base of data center equipment. This method avoids overestimation that can be caused by tracking data center load for projects that have not yet selected a power provider, but requires many inputs and assumptions developed from limited publicly available data, proprietary market analyst data, and review by industry representatives and stakeholders. The lack of direct energy data available in a sector with rapidly evolving technologies limits the analysis in this report, especially when trying to understand and estimate future energy demand scenarios.

The results presented here indicate that the electricity consumption of U.S. data centers is currently growing at an accelerating rate. Figure ES-1 shows a compound annual growth rate of approximately 7% from 2014 to 2018, increasing to 18% between 2018 and 2023, and then ranging from 13% to 27% between 2023 and 2028. This surge in data center electricity demand, however, should be understood in the context of the much larger electricity demand that is expected to occur over the next few decades from a combination of electric vehicle adoption, onshoring of manufacturing, hydrogen utilization, and the electrification of industry and buildings. Research initiatives are needed not merely to identify strategies to meet data centers' future energy needs, but also to help stakeholders use this relatively near-term electricity demand for data centers as an opportunity to develop the leadership and strategic foundation for an economy-wide expansion of electricity infrastructure.

Aqueous Film-Forming Foam (AFFF)

1 Introduction

Aqueous film-forming foam (AFFF) is a highly effective firefighting product intended for fighting high-hazard flammable liquid fires. AFFF products are synthesized by combining hydrocarbon foaming agents with fluorinated surfactants to achieve a product that has been used at military installations, civilian airports, petroleum refineries, bulk storage facilities, and chemical manufacturing plants (Hu et al. 2016; CONCAWE 2016).

This fact sheet targets local, state, and federal regulators and tribes in environmental, health and safety roles, as well as AFFF users at municipalities, airports, and industrial facilities, and is not intended to replace manufacturer specifications or industry guidance for AFFF use. The information provided is a high-level summary on AFFF use, the associated hazards, and how to reduce and eliminate potential harm to human health and the environment. Additional information is available in the Guidance Document.

ITRC has developed a series of fact sheets that summarizes recent science and emerging technologies regarding PFAS. The information in this and other PFAS fact sheets is more fully described in the ***ITRC PFAS Technical and Regulatory Guidance Document (Guidance Document)*** (*https://pfas-1.itrcweb.org/*).

This fact sheet outlines methods to properly identify, handle, store, capture, collect, manage, and dispose of AFFF to limit potential environmental impacts, and includes:

- Definition of AFFF
- Best Management Practices for AFFF use
- Regulations Affecting Sale and Use
- Foam Research and Development

2 What is AFFF?

Class B firefighting foams are commercial surfactant solutions that are designed and used to combat Class B flammable fuel fires. For the purpose of this fact sheet, Class B foams can be divided into two broad categories: fluorinated foams that contain PFAS and fluorine-free foams (F3) that do not contain PFAS.

There are six groups of Class B foams that contain PFAS and four groups of Class B foams that do not. Figure 1 illustrates all categories of Class B foams. This fact sheet addresses only AFFF because it is the most widely used and available type of Class B foam.

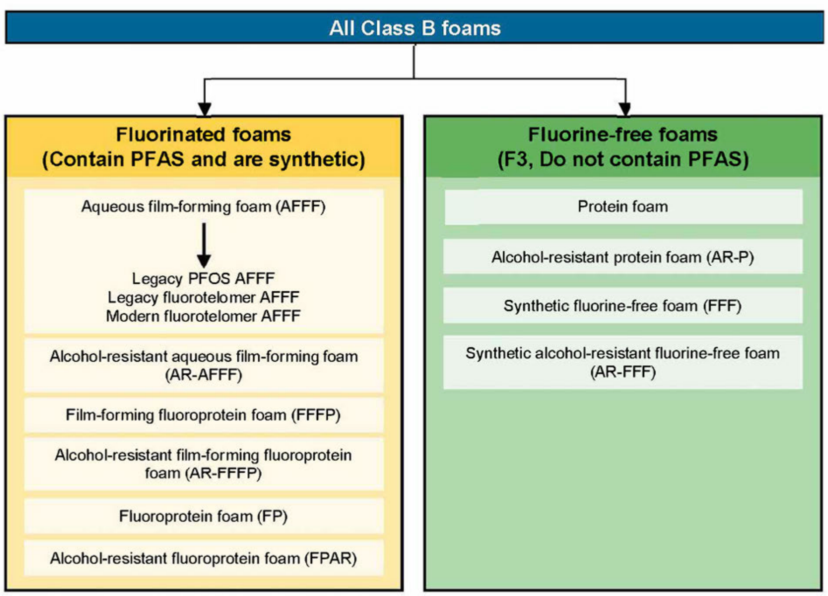

Figure 1. Types of Class B foams.
Source: S. Thomas, Wood, PLC. Used with permission. PFAS-1, Figure 3-2.

AFFF is considered a fluorinated foam and when mixed with water, the resulting solution achieves the interfacial tension characteristics needed to produce an aqueous film that spreads across the surface of a hydrocarbon fuel (petroleum greases, tars, oils, and gasoline; and solvents and alcohols) to extinguish the fire and to form a vapor barrier between the fuel and atmospheric oxygen to prevent re-ignition. This film formation is the defining feature of AFFF.

AFFF has been used at chemical plants, flammable liquid storage and processing facilities, merchant operations (oil tankers, offshore platforms), municipal services (fire departments, firefighting training centers), oil refineries, terminals, and bulk fuel storage farms, aviation operations (aircraft rescue and firefighting, hangars), in some industrial fire extinguishers, and military facilities.

There are three types of AFFF. Each is presented in Figure 1:

- legacy PFOS AFFF (manufactured in the US from the late1960s through 2002)
- legacy fluorotelomer AFFF (contain some long-chain PFAS) (manufactured in the US from the 1970s until 2016)
- modern fluorotelomer AFFF (short-chain PFAS became the predominant fluorochemicals used in manufacturing in response to USEPA 2010/2015 voluntary PFOA Stewardship Program)

Most foam manufacturers now produce Class B F3s, and evaluation of the performance of these foams is an important consideration for future purchase decisions. As part of preplanning for replacement foams, it is important to ensure that the Class B F3 is not considered a regrettable substitution over AFFF and can achieve the required performance specifications for the target flammable liquid hazards (FFFC 2016). The 2022 National Defense Authorization Act (NDAA) requires a new Military Specification (Mil-Spec) for PFAS-free foams by January 2023. As of May 2022, F3s do not meet the performance requirements of the current Mil-Spec and are not used at federal- and FAA-regulated facilities (FAA 2021).

The NDAA of fiscal year 2020 (signed into law Dec 20, 2019) requires the DOD to phase out its use of AFFF at all military installations by Oct. 1, 2024, with limited exceptions, and immediately stop military training exercises with AFFF. The secretary of the Navy must publish specifications for PFAS-free firefighting foam at all military installations and ensure that the foam is available for use by Oct. 1, 2023. The NDAA of fiscal year 2022 also addresses AFFF, specifically requiring new reviews and guidance to prevent and mitigate AFFF spills. In October 2021, the USEPA published the PFAS Strategic Roadmap: EPA’s Commitments to Action 2021–2024 (USEPA 2021 Ref#2223). The USEPA’s stated goals for addressing PFAS are focusing on research, restriction, and remediation. The strategic roadmap includes actions across the different divisions of USEPA. More information about USEPA’s actions in 2021 to address PFAS are available on their website (USEPA 2021 Ref#2223).

3 Best Management Practices (BMPs) for Class B AFFF Use

Firefighting foams are an important tool to protect human health and property from flammable liquid fire threats. Proper management and usage strategies combined with the ongoing refinement of environmental regulations will allow an informed selection of the viable options to sustainably use firefighting foams. BMPs should be established for the use of any firefighting foam to prevent possible releases to the environment that can lead to soil, groundwater, surface water, and potentially drinking water contamination. The discharge of firefighting foams to the environment is of concern because of the potential negative impact they can have on ecosystems and biota. AFFF, due to the presence of PFAS, poses a unique challenge to protecting the environment when it is released. Specifically, for AFFF, the amount of PFAS from foam that may enter groundwater depends on the type and amount of foam used, the degree of containment, when and where the foam was used, the type of soil and the depth to groundwater. AFFF is typically discharged on land but can run off into surface water or stormwater or infiltrate to groundwater. A conceptual site model (CSM) is presented in Figure 2.

BMPs start with pre-planning and deciding which foam to keep in stock. The team should consider key factors such as these:

- Whether F3 alternatives can meet site-specific performance requirements
- Site-specific evaluation of likely fire hazards and potential risks for life, public safety, and property
- Potential environmental, human health, and financial liabilities associated with AFFF releases
- Site constraints, including existing equipment retrofit requirements to adapt to alternate foams

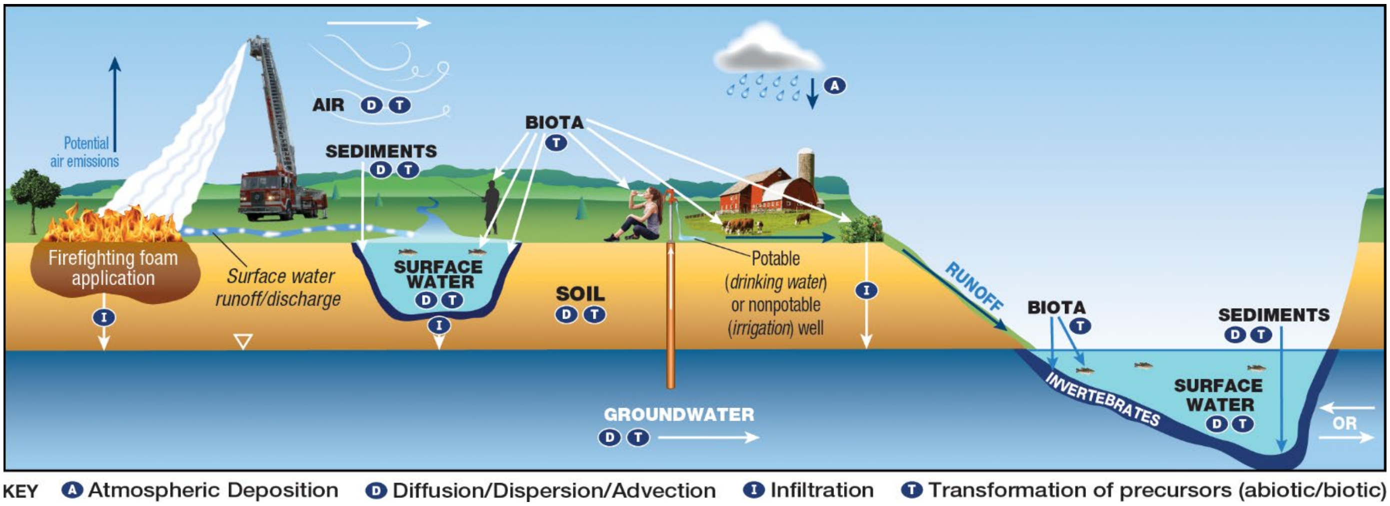

Figure 2. CSM for fire training areas.

Source: Adapted from figure by L. Trozzolo, TRC. Used with permission. PFAS-1, Figure 2-19.

BMPs should consider the entire life cycle for AFFF, Figure 3, including procurement and inventory, foam systems and operations, emergency firefighting operations, immediate investigative and clean-up actions, treatment and disposal and system replacement.

The procurement and inventory of foam should be carefully considered. Foams should be selected that meet the performance specification requirements governing the use. Foams procured should be documented, labelled clearly and adequately contained. Foam use and disposal should be carefully tracked and recorded.

When evaluating foam systems and operations, from fixed-system testing, mobile firefighting equipment testing and appropriate training exercises, engineering and administrative controls as well as personal protective equipment (PPE) should be carefully evaluated. During emergency firefighting operations following a release of firefighting foam, PPE should be used correctly, maintained, and decontaminated routinely. Immediate investigative and clean-up actions include initial mitigation efforts such as source control, containment tactics, and recovery tactics.

The treatment and disposal of AFFF products and environmental media impacted with PFAS can be complex, time consuming, and costly. Practitioners should be aware of approved and available disposal options prior to the generation of PFAS-impacted waste or the start of an AFFF replacement project to avoid potentially lengthy waste storage timeframes. Currently, available disposal options for AFFF and PFAS-impacted materials are limited and each option has its advantages and disadvantages. More information is in the Guidance Document, as well as in the PFAS Regulatory Programs Summary Table (see the External Data Tables on https://pfas-1.itrcweb.org).

Firefighting foam replacement is complex and could require a complete system review and, potentially, redesign and modification of system components to meet the new objectives or material and performance requirements. Foam replacement should include an evaluation of specific hazards and application objectives, a review of applicable performance standards, an understanding of engineering requirements for foam product storage and application, and a check to ensure that the foam product is approved for use for the specific hazards being mitigated.

Figure 3. Life cycle considerations for AFFF.

Source: S. Thomas, Wood, PLC. Used with permission. PFAS-1, Figure 3-1.

4 Regulations Affecting the Sale and Use of AFFF

There are many State, Federal, and International regulations and guidance documents governing the procurement, use, and disposal of AFFF. Activities range from AFFF take-back programs and prohibition of manufacture, sale, use, and import of AFFF through to restrictions and requirements for disposal. More information is in the Guidance Document, as well as in the PFAS Regulatory Programs Summary Table (see the External Data Tables on https://pfas-1.itrcweb.org).

5 Foam Research and Development

A substantial amount of research related to AFFF alternatives and replacement chemistries has recently been completed or is being considered at the time of publication. Several organizations globally have made investments in research and development around AFFF from the assessment of their use, environmental impacts, as well as socioeconomic impacts of transiton to and performance specifications of F3 alternatives. For more information related to this topic, please refer to the Guidance Document.

6 References and Acronyms

The references cited in this fact sheet and further references can be found at https://pfas-1.itrcweb.org/references/. Reference numbers are included in this fact sheet for non-unique citations in the Guidance Document reference list.

The acronyms used in this fact sheet and in the Guidance Document can be found at https://pfas-1.itrcweb.org/acronyms/.

Per- and Polyfluoroalkyl Substances (PFAS) Team Contacts

Sandra Goodrow • New Jersey Department of Environmental Protection
609-940-4164 • Sandra.Goodrow@dep.nj.gov

Kate Emma Schlosser • New Hampshire Department of Environmental Services
603-271-2910 • KateEmma.A.Schlosser@des.nh.gov

July 2022

U.S. National Science Foundation

NSF ENGINES INAUGURAL AWARD

GREAT LAKES ReNEW
GREAT LAKES WATER INNOVATION ENGINE
Current

Building a Circular Water Economy in the Great Lakes

Conventional water and wastewater treatment flushes valuable resources down the drain and leaves harmful contaminants behind. At a time of growing water scarcity, **Great Lakes ReNEW** will transform waste into wealth for the Great Lakes Region. We will leverage our abundant share of global freshwater to drive innovation, safeguard our domestic water resources and supply chains, attract and retain water intensive industries, and position the U.S. as a global leader in economic growth and emerging industries.

In January 2024, the National Science Foundation awarded up to $160 million over 10 years to Current, a Chicago-based water innovation hub, to establish, develop, and grow the Great Lakes Water Innovation Engine: Great Lakes ReNEW.

ReNEW's goal is to accelerate the transition to a circular blue economy through:

- **Use-inspired R&D** on selective separation and resource recovery of nickel, cobalt, lithium, nitrogen and phosphorus, and elimination of PFAS and other contaminants from water and wastewater.
- **Translation of innovation to market**, connecting the assets of leading regional water hubs and testbeds, launching and investing in dozens of watertech startups, and building a centralized testbed for their products.
- **Workforce development** to train people for quality jobs and careers and support K-12 STEM education.

ReNEW's **economic impact** will create new industries and jobs in the Great Lakes region, positioning the region as a leader in water innovation and circular economy practices. ReNEW's **environmental impact** will improve human and planetary health through the discovery, development, and deployment of new technologies and practices for circular water management.

Highlights

- **Funding:** Up to $160 million support from the National Science Foundation's Regional Innovation Engines program, with additional public (state and local), private (corporate), and philanthropic investments.
- **Leadership:** Led by Current's CEO and Engine Principal Investigator Alaina Harkness; Co-PI is Dr. Junhong Chen of Argonne National Laboratory and the University of Chicago's Pritzker School of Molecular Engineering.

- **Regional Scope:** Core partners span six Great Lakes states—Illinois, Indiana, Michigan, Minnesota, Ohio, and Wisconsin.
- **Collaboration:** 50+ partners including research institutions, industry leaders, investors, government entities, and nonprofits.

Objectives

- **Economic Growth:** Create new opportunities; support emerging industries like AI, quantum, semiconductor manufacturing, etc; strengthen supply chains; and spur job creation in the region.
- **Resource Recovery:** Innovate to extract and reuse energy, nutrients, and materials vital to the clean energy transition (lithium, nickel, cobalt) while removing contaminants (PFAS) from wastewater.
- **Sustainability:** Shift from linear to circular water use; tackle water scarcity and quality issues by removing harmful chemicals and contaminants; improve the energy efficiency of water treatment technologies.
- **Workforce Development:** Create career pathways and training programs that strengthen the water workforce through industry, community, and educational partnerships.

Building on Success

Since 2016, Current has put water innovation front and center as an economic opportunity and solution to managing growing global water risk. Current has raised more than $58 million to support water innovation and economic development, launched 11 technology pilots, supported more than 40 water tech startups, and educated more than 40,000 people through events and convenings.

Join us in transforming waste into wealth and health and driving the future of water innovation in the Great Lakes region. For more information and to explore partnership opportunities, visit currentwater.org & greatlakesrenew.org

CLEVELAND WATER ALLIANCE

About Us Programs Engage News & Media

CHANGING HOW WE VALUE WATER

The need has never been so urgent: innovation around how we use, protect, remediate, and fortify water resources and infrastructure is critical to the future of our industries, communities, and economies - here in Ohio and around the world.

+ Join us at the leading edge

WHAT IS CWA?

UNIVERSITY OF ILLINOIS CHICAGO

The Freshwater Lab

WATER REUSE SOLUTIONS FOR DATA CENTERS

WHAT ARE DATA CENTERS AND WHY ARE THEY BECOMING MORE PREVALENT?

Data centers are physical facilities that house computers, servers, and other equipment used to store, manage, and process vast amounts of data for websites, apps, cloud computing, digital services, and AI. Their prevalence is increasing with the rise of AI and other fast, interactive computing technologies, driving the need for ever-larger facilities to process data more efficiently. Speed, size, and demand are key factors in this expansion.

As of 2024, there are at least 8,726 data centers worldwide. In the U.S., 3,225 data centers are distributed across various cities, including 145 in Chicago, 68 in New York City, and 107 in Cleveland. Canada has 251 data centers, with Toronto leading at 84, followed by Montreal with 54 and Vancouver with 26.[1]

Data Centers in the Midwest

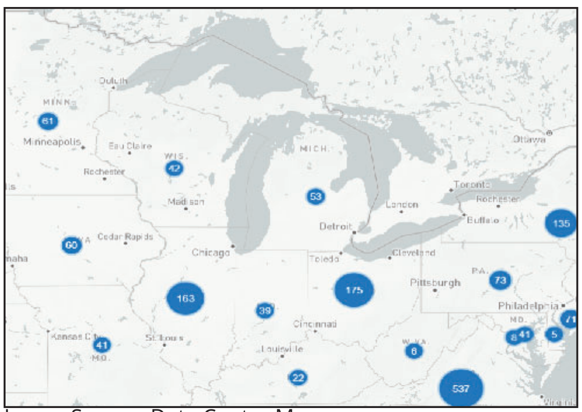

Image Source: Data Center Map

WHY DO DATA CENTERS NEED WATER?

Most data centers rely on municipal drinking water, using the same high-quality supply as residents at the same cost, primarily to cool their computers, servers, equipment, and other systems. New facilities supporting AI can use millions of gallons daily—about the same as a small town.

There are two main cooling methods: air conditioning powered by electricity or water-based systems. Evaporative cooling, while energy-efficient, consumes large amounts of water, straining resources in water-scarce areas. This makes balancing water conservation and energy efficiency difficult. Even electricity-based cooling has a hidden water cost due to power generation.[2] Ultimately, there's no way around high water use in data center cooling.

Using renewable energy can reduce both carbon and water footprints, while recycled water or closed-loop systems help cut water intake, improving efficiency and cost-effectiveness.

HOW MUCH WATER DO THEY NEED?

Google's hyperscale data centers, which power services like Gmail and Google Drive, average 550,000 gallons per day, while smaller centers use about 18,000 gallons. In the U.S., where individuals use 132 gallons daily, a large data center consumes as much water as 4,200 people, ranking data centers among the top 10 industrial water users.[3]

With over 3225 data centers in the U.S., about 20% were drawing water from stressed western watersheds by late 2021, worsening regional shortages.[4]

Meanwhile, the Midwest is emerging as a hyperscale data center hub, driven by AI and cloud demand:

- Chicago, IL: Data centers are expected to grow from 66 to 122 by 2028, with capacity rising from 997 MW to 1,839 MW. [5]
- Columbus, OH: The number of data centers is set to more than double from 40 to 85 by 2028.[6]

Globally, AI's annual water consumption is projected to reach 4.2 to 6.6 billion cubic meters by 2027—four to six times Denmark's annual usage.[7]

Comparative Water Impact

Data Center Capacity or Classification	Annual Water Consumption	Perspective
1 MW	6.7 Million Gallons	Annual water consumption of 225 Americans
Hyperscale	200 million gallons	Annual water consumtion of 6700 Americans

Source: Accenture and dgtlinfra

SUSTAINABLE SOLUTIONS

By adopting water recycling and advanced treatment systems, Great Lake states can promote economic growth while preserving precious freshwater. This new approach can position the region as the world's Water Belt able to absorb new industry, agriculture, and population.

Water reuse systems are already in place across the U.S. and the world. States like California and Texas are integrating recycled water in their drinking water supplies, Arizona reuses water for manufacturing and energy production, and Virginia protects military assets by capturing and reusing stormwater. Drought and disaster have brought these states to swiftly implement water reuse.[8] Bringing water recycling into use now can avert such outcomes in Great Lake states.

Great Lakes communities can produce a range of high quality water to accommodate different uses. By collaborating on fit-for-purpose systems, utilities and industries can together advance technology in the region. Greater water availability brings more revenue and secures the potential for future growth.

Data centers can optimize water use with smart systems, continuously recycling water or producing it for external use, avoiding strain on local sewer systems.

Water reuse options for data centers include:

- Closed-loop systems with onsite water recycling, including rainwater harvesting.
- Municipalities supplying data centers with recycled water.
- Using data center water discharge for agricultural or industrial purposes.
- Data center partnerships for watershed replenishment or aquifer recharge.

Closed-Loop Systems

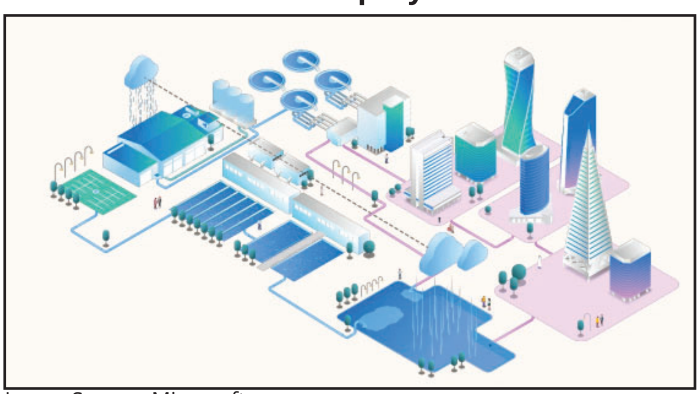

Image Source: Microsoft

WATER RECYCLING TEAM

The Water Recycling Team at the University of Illinois Chicago looks forward to working with Great Lakes communities on the design and implementation of water recycling systems.

> Rachel Havrelock
Director, The Freshwater Lab, UIC
Email: raheleh@uic.edu

> Krishna Reddy
Director, The Sustainable Engineering Research Lab, UIC
Email: kreddy@uic.edu

> Teresa Cordova
Director, UIC Great Cities Institute
Email: tcordova@uic.edu

> Ahmed Rachid El-Khattabi
Interim Director, Environmental Finance Center, UNC Chapel Hill
Email: arelkhattabi@gmail.com

> Matthew Wilson
Associate Director of Economic and Workforce Development, UIC Great Cities Institute
Email: mwilso25@uic.edu

ENDNOTES

1. Data Center Map. (n.d.). Data Center Map - Colocation, Cloud and Connectivity. Retrieved February 17, 2025, from https://www.datacentermap.com/
2. Higgins, A. (2024, November 13). What is water usage effectiveness (WUE) in data centers? Interconnections - The Equinix Blog. https://blog.equinix.com/blog/2024/11/13/what-is-water-usage-effectiveness-wue-in-data-centers/
3. Pinheiro Privette, A. (2024, October 11). AI's challenging waters. Center for Secur Water- Department of Civil and Environmental Engineering. https://cee.illinois.edu/news/AIs-Challenging-Waters
4. Ibid.
5. Davis, J. (2024, November 18). Big data centers, big rewards for states? The Council of State Governments - Midwest. https://csgmidwest.org/2024/11/18/big-data-centers-big-rewards/
6. Ibid.
7. Center for Secure Water- Department of Civil and Environmental Engineering. (2024, October 11).
8. Swinhoe, D. (2023, November 23). AWS using reclaimed wastewater for data center cooling at 20 locations. Data Center Dynamics. https://www.datacenterdynamics.com/en/news/aws-using-reclaimed-wastewater-for-data-center-cooling-at-20-locations/