

For professional investors – Marketing communication – April 2026

WATER: THE HIDDEN CONSTRAINT IN THE DIGITAL ECONOMY



BNP PARIBAS
ASSET MANAGEMENT

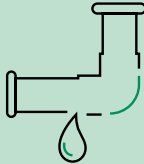


The sustainable
investor for a
changing world

INTRODUCTION

Structural megatrends are intensifying pressure on water resources, increasing scarcity and posing growing challenges to long-term sustainability and economic growth. In response, the Environmental Strategies Group invests in companies providing solutions across the water value chain, from upstream activities such as procurement & equipment, through mid-stream production & distribution, to downstream operation & end use.

In this paper, we focus on downstream water use, with particular attention to the relationship between artificial intelligence (AI) and water, to demonstrate why water is becoming a driver of both investment risk and opportunity.

Our Water Framework

 Water Equipment & Components	 Water Utilities & Infrastructure	 Water Efficiency & Technologies
Industrial & Municipal Water Equipment Building Water Equipment Water Irrigation Equipment Water Specialty Components Upstream: Procurement & Manufacturing	Water Infrastructure & Concessions Water Utilities – Europe Water Utilities - US Water Utilities - Rest of World Mid-stream: Production & Distribution	Water Treatment & Technologies Water Testing & Quality Water Solutions Water Monitoring & Efficiency Downstream: Operation & Use

Forecasts for new AI hyperscale data centres are growing exponentially and so too is energy demand, for both behind-the-meter and grid connected solutions. Subsequently, AI and data centres are the single biggest driver behind additional US electricity needs to 2033 (Source: BNP Paribas Exane). A new report estimated that global AI and data centres could use as much energy as Japan by 2030.¹

¹ [AI is set to drive surging electricity demand from data centres while offering the potential to transform how the energy sector works - News - IEA \(2025\)](#)

Whilst energy remains a critical challenge, another is emerging as the supply chain supporting AI demand evolves and expands: water. Hyperscale facilities can require more than 2 billion litres annually and data centres more generally in the United States alone, consumed an estimated 66 billion litres of water in 2023. The latter amount is comparable to the annual domestic water use of a midsized city.² Furthermore, a recent survey conducted by Morgan Stanley found that 72% of respondents cite water availability as a top overall concern and a cause of slowing data centre projects.³

Microchips or “chips”

Chips, used for computing in data centres, are produced from semiconductors, which have an even higher water footprint. Semiconductors, manufactured in semiconductor fabrication plant or “fabs”, are the foundational hardware that powers AI, whilst AI is transforming the way they are designed, manufactured and optimised. Chip manufacturing relies heavily on ultrapure water, with a single fabrication plant capable of using around 14 billion litres per year.⁴ For every unit of ultrapure water required, additional municipal water is needed in the production process. At a global level, total water consumption by the semiconductor industry is estimated to be comparable to that of a city with a population of approximately 7.5 million people.⁵ These amounts are expected to more than double from 2024 to 2028

This growing reliance on water exposes operating companies to a range of **risks**, including transition, regulatory and reputational risks, particularly in water stressed regions. At the same time, the expansion of AI and digital infrastructure is creating **opportunities**. These include reduced operating costs (1) through improved water efficiency and lower dependence on scarce natural resources, (2) for companies that enable water related technologies and solutions across the value chain, and (3) for asset owners to allocate capital to scalable water solutions that support the long-term resilience of the technology sector. These dynamics are reinforced by a range of structural tailwinds, including evolving regulation, permitting requirements, water positive corporate targets, and improving methodologies for measuring and disclosing water use in data centres.

This paper dives into two factors, which are framed through **direct uses of water** (on-site data centre cooling) and **indirect uses of water** (upstream devices including semiconductor manufacturing).⁶ Water is also included in the AI supply chain through power generation as global fossil fuel water use has risen sharply. Coal dominates global water impacts, accounting for 62% of withdrawals and 57% of consumption, while crude oil represents 34% of consumption, natural gas 33% of withdrawals, and total fossil fuel wastewater reached 43×10^9 m³ in 2019, mainly from coal (52%) and crude oil (46%).⁷ This latter point, however, is omitted from this paper as it is mainly addressed by switching to renewables as wind and solar PV do not require cooling systems and have the least water usage compared to thermal power generation.⁸

2 TNFD (2026), Nature-related Issues in the Technology Sector: Dependence on water by the semiconductor and data centre industries.

3 Morgan Stanley (2025), AI's Growing Thirst for Water.

4 TNFD (2026), Nature-related Issues in the Technology Sector: Dependence on water by the semiconductor and data centre industries.

5 TNFD (2026), Nature-related Issues in the Technology Sector: Dependence on water by the semiconductor and data centre industries.

6 Morgan Stanley (2025), AI's Growing Thirst for Water.

7 [The global water footprint of fossil fuels: The role of the energy-water quality nexus - Astrophysics Data System](#) (2021)

8 hydroelectric water use is not considered consumption.
Source: United States Data Center Energy Usage Report

DIRECT WATER USAGE: DATA CENTRES

Predictions for data centre growth out to 2040 are staggering. It is estimated that:

- 10× more compute power will be required for AI workloads compared with traditional data centre applications (source: Goldman Sachs, 2024)
- There are ~3,000 data centres currently under construction or planned globally, expected to be completed by 2030 (source: IRMI, 2026)
- There is USD 7 trillion in total global data centre investment projected through 2030
- USD 77.7 billion in US data centre construction started in 2025 alone, representing 190% year on year growth (source: ConstructConnect, 2025)

This **new capacity** alongside **increasing processing capacity** creates significant heat. Data centres have been shown to raise temperatures around 2 degrees Celsius in their local area. This is exacerbated by global warming and increasing frequency and severity of heatwaves (in 2022, summer heat waves in the UK and US caused several data centres to go offline). As such, operators need to find a way to consistently and reliably cool AI infrastructure. This can be done by several cooling technologies.

Cooling is a critical function within data centres, as the significant heat must be removed to maintain performance due to thermal throttling, reliability and equipment lifespan. Cooling strategies operate at two levels: server-level cooling, which removes heat directly from hardware components, and facility-level cooling, which expels that heat from the building.

Server Level Cooling

Server-level cooling focuses on extracting heat as close to the source as possible, namely from servers and their internal components (i.e., the GPU and CPU). To date, this has mostly been achieved through air cooling, where fans and air-conditioning systems circulate air through server racks. Hot air generated by servers is expelled, while cooler air is drawn in, often using rack-mounted fans and hot-aisle/cold-aisle containment layouts. Air cooling is simple, well-understood and relatively low-cost to deploy. However, given the recent trend towards increased rack compute densities there is an increasing trend toward higher cooling requirements per rack which air cooling cannot facilitate given its lower heat capacity compared to liquid cooling alternative.

As compute densities rise, liquid cooling is becoming the dominant server-level solution. Liquid cooling circulates a refrigerant or water-based coolant close to, or directly over, heat-generating components, allowing heat to be absorbed and removed far more efficiently than with air. Direct-to-chip cooling uses cold plates attached to CPUs and GPUs⁹ to remove heat at the source, while immersion cooling submerges entire servers in a non-conductive liquid that transfers heat directly to the surrounding fluid. These approaches offer significantly higher efficiency, support dense AI workloads and reduce the need for large airflow systems, though they require more complex design and specialised infrastructure.

9 A CPU (Central Processing Unit) and a GPU (Graphics Processing Unit) are both critical computing components, but they are architected for different purposes and excel in distinct workloads.

Source: IRMI, 2026

Source: New Scientist, March 2026)

Facilities Level Cooling

Once heat has been transferred away from servers, facility-level cooling systems remove it from the building. One common approach is air-based heat rejection, which relies on chillers, cooling towers and HVAC systems to transfer heat to air and expel it outside. While widely deployed, these systems are highly energy-intensive and contribute materially to a data centre's electricity demand.

One example of facility-level cooling is evaporative cooling, which uses water to absorb heat and release it through evaporation. In this process, warmed water or refrigerant passes through a cooling system where evaporation carries heat into the atmosphere. Evaporative cooling is more energy-efficient than purely air-based systems and can be particularly effective in dry climates, but it increases water consumption and raises sustainability concerns as AI workloads scale.

Some data centres also deploy free cooling, which leverages favourable external conditions such as cool ambient air or naturally cold-water sources to minimise mechanical cooling. Air-side economisation draws in cool outside air directly, while water-side economisation uses naturally cool water for heat exchange. These approaches are highly energy-efficient but are inherently climate-dependent and less viable in hot or humid regions.

In practice, many modern data centres use hybrid cooling systems that combine air, liquid, evaporative and free-cooling techniques.¹⁰ These integrated designs allow operators to balance energy efficiency, water use and performance requirements, optimising cooling strategies as workloads and environmental conditions evolve.

INDIRECT WATER USAGE: SEMICONDUCTORS

Capital expenditure for a new fabrication facility can reach up to USD 20 billion, depending on scale and complexity. Of this total, up to approximately 10% is typically allocated to water-related infrastructure, including ultra-pure water (UPW) production and advanced wastewater treatment systems. At this scale of investment, water clearly emerges as one of the most critical inputs to semiconductor manufacturing.¹¹

How does the industry use water?

The role of water in the semiconductor manufacturing process is highly specialised. Ultra pure water is processed to achieve extremely stringent specifications, including very high electrical resistivity, ultra low organic content, minimal particle presence and near sterile conditions. These purity levels exceed those of standard purified water, requiring the removal of ionic contaminants, dissolved gases, organic compounds and particulates to parts per billion concentrations.

Within semiconductor fabrication, ultra pure water serves several essential functions. It is used extensively to clean silicon wafers by removing particles and chemical residues; to act as a contaminant free base for the preparation of chemicals used in etching and photolithography; to rinse wafers following etching and chemical mechanical planarization processes; and to support precise thermal management of photolithography equipment, where temperature stability is critical to manufacturing accuracy and yield.

10 [How Hybrid Design Could Be the Future of Data Centres | Data Centre Magazine](#) (2025)

11 [semiconductors-the-water-challenge-us.pdf](#) (2025)

How much water is used?

Semiconductor fabrication is highly water-intensive, reflecting the critical role of ultra-pure water in advanced manufacturing processes. A typical semiconductor fabrication plant consumes up to five million gallons of municipal water per day (~10 to 16 million bottles) to produce the volumes of ultra-pure water required for production, with the most advanced facilities using around 4.5–7 litres of ultra-pure water per square centimetre of processed wafer. This process is inherently resource-intensive: producing 1,000 gallons of ultra-pure water typically requires between 1,400 and 1,600 gallons of municipal water input.¹²

FIGURE 1. Water is required in volume across the semiconductor manufacturing process...

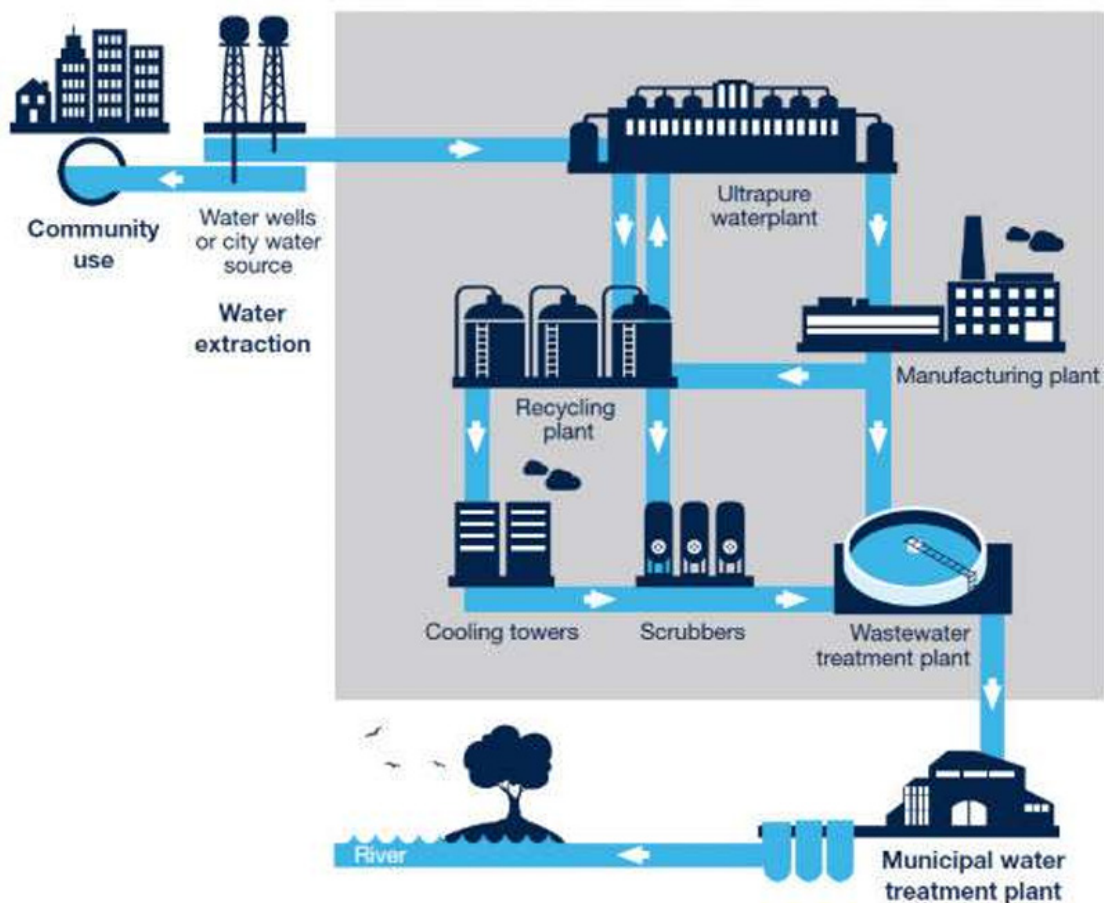


Image Source: Barclays (2026). Semis' thirst for water: A shifting global picture

While water recycling is increasingly embedded within fab operations, average recycling rates currently range from 65–75%, with next-generation facilities targeting significantly higher reuse rates of 85–90%.¹³ Despite these efficiency gains, water remains a meaningful operating cost, with annual water-related expenses typically ranging from USD \$2-8 million per fabrication plant, depending on production scale and local water pricing.

12 Morgan Stanley (2025), AI's Growing Thirst for Water.

13 [Ultrapure Water Systems in Semiconductor Manufacturing Explained | AXEON Water](#) (Date:NA)

PFAS Risks in Semiconductor manufacturing

The semiconductor industry is closely linked to PFAS (per- and polyfluoroalkyl substances) because these chemicals play a critical functional role in advanced chip manufacturing, even as they pose growing environmental and regulatory concerns. PFAS are used in semiconductor fabrication primarily for their chemical stability, resistance to heat, and surfactant properties, which are essential in processes such as photolithography, etching, and cleaning. For example, they're used in photoresists, developers, anti-reflective coatings, and specialty wet-process chemicals that must perform reliably under extreme conditions and at nanometre scales.

However, these same properties make PFAS persistent in the environment, meaning they can accumulate in water, soil, and human bodies if released through wastewater, air emissions, or improper waste handling. As fabs are among the most water-intensive industrial facilities, even trace PFAS discharges can become a community concern, particularly where wastewater treatment systems are not designed to fully remove them. This has drawn increased scrutiny from regulators and local stakeholders, placing pressure on semiconductor companies to invest in advanced water treatment, closed-loop systems, substitution research, and tighter supply-chain controls. In short, PFAS are used because they enable the performance and yield required for modern chips, but their environmental persistence is forcing the industry to rethink materials, water management, and accountability as fabs expand to meet AI-driven demand.

Source: [FINAL-PFAS-Consortium-Background-Paper.pdf](#)

Geography matters

The location of hyperscalers is influenced by a range of factors, including: (1) high-quality internet connectivity to minimise latency given the significant volumes of downstream data; (2) access to low-cost and readily available power; and (3) permitting, planning, and regulatory considerations, among others.

However, these factors are not necessarily aligned with natural resource availability. As a matter of fact, the global distribution of data centres has, in practice, increasingly overlapped with regions that are already experiencing water stress. One contributing factor may be that locations well suited to large-scale solar power generation are often characterised by limited water availability.

The global semiconductor manufacturing footprint reflects a combination of structural industry consolidation, rising capital intensity and increasingly complex site-selection requirements. As manufacturing needs evolve alongside advances in technology, pressure on existing fabrication capacity has intensified. Limited availability of operational fabs and cleanroom space has driven a wave of new capital projects and expansions, pushing the value of existing facilities to a premium. At the same time, companies are increasingly evaluating greenfield developments, often supported by local and national government incentives aimed at strengthening domestic semiconductor supply chains.

These expansion efforts are occurring against a backdrop of persistent operational and strategic constraints. There remains a shortage of both new and second-hand production tools, with lead times from original equipment manufacturers frequently exceeding one year, constraining the pace at which new capacity can be brought online. Decisions around fab location must therefore balance a range of internal and external considerations. Internal factors include the scale of financial investment required, existing customer loading and supply agreements, and broader risk-mitigation strategies. Resource availability has become central here, with reliable access to high-quality water now a key determinant of site viability.

This has increased the importance of diversified and resilient water sourcing arrangements, alongside geographic diversification to reduce exposure to disruption.¹⁴

In parallel, external factors including geopolitical dynamics such as trade policy and tariffs, global competitive positioning, intellectual property protection, proximity to local supply chains, access to skilled labour and the availability of government incentives all play a critical role in shaping fab location decisions. Together, these factors have driven the current global distribution of semiconductor fabrication facilities, illustrated by the existing landscape of fabs worldwide. As we can see below, many of these Fabs exist in medium to high water stress areas.

Footprint of global fabs, combined with a water risk map

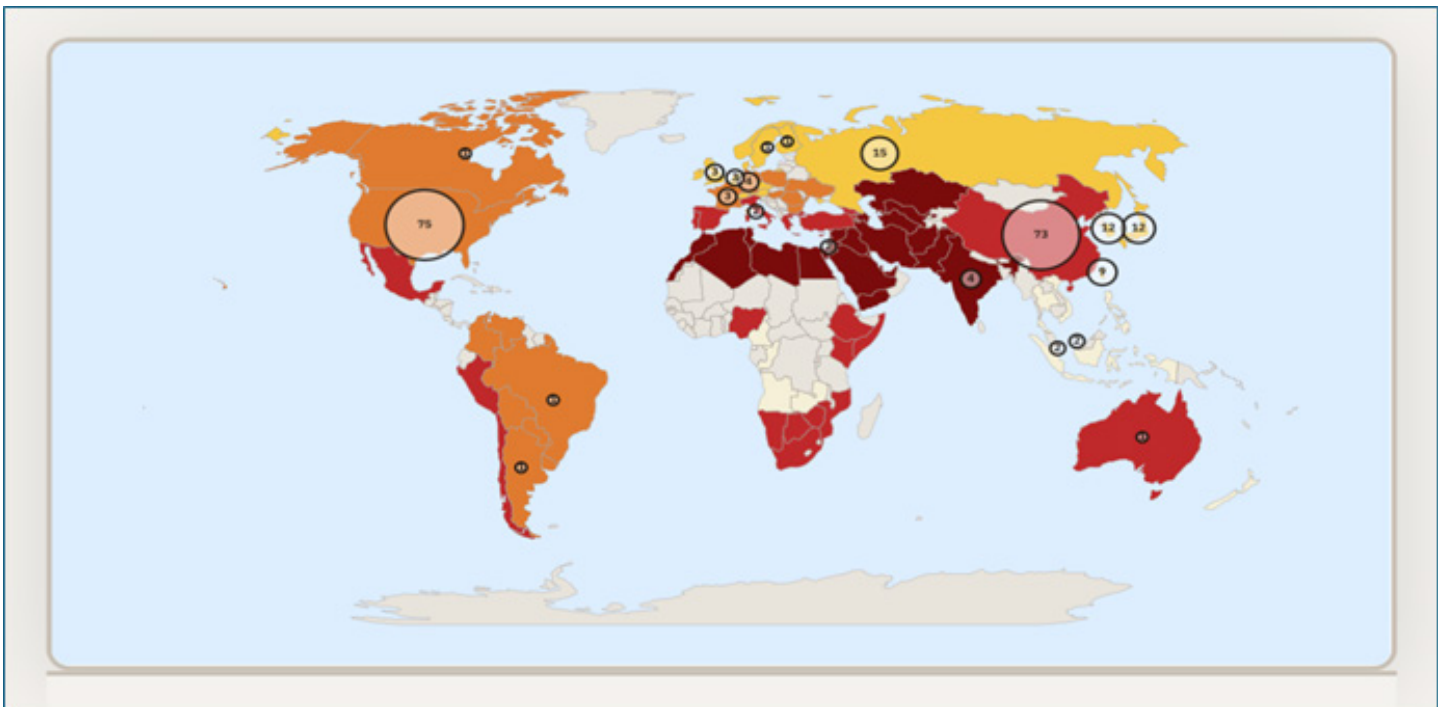


Image Source: AI generated image overlaying a global water stress map and the global presence of semiconductor fabs. Original images are sourced from Technology of Global Affairs and the World Resources Institute
 *Increasing size of circle indicates larger presence of data centres, ** darker red indicates increasing water stress

14 [Semiconductor manufacturing strategy: Where in the world to locate a fab or cleanroom? | International Symposium on Microelectronics \(2019\)](#)

To further illustrate through specific case studies¹⁵:



The US share of global semi production expected to triple between 2022 and 2032, supported by new fabs in Arizona, Idaho, Ohio and Texas. Arizona is a major hub despite being a highly water-stressed region, with TSMC alone committing USD 165 billion to five fabs. In 2023, water shortages forced Arizona to cut Colorado River extraction by 592,000 acre-feet -21% of the river's allocation and 9% of state water use-raising concerns about competition with housing and agriculture.



Taiwan is central to the global semiconductor supply chain, with TSMC accounting for over 90% of advanced microchip revenues. Taiwan relies on typhoons for around 70% of its rainfall, and the absence of typhoons during the 2020-21 season triggered a severe drought, with reservoir levels falling below 5% capacity. TSMC spent NT 800 million (USD 28.6 million) on water trucking-60% more than its emergency budget.



Since 2019, Chinese semiconductor firms have invested more than USD 63 billion across 73 fabs. Many facilities are located in water-stressed and polluted regions, where the semiconductor sector accounts for approximately 27% of total industrial water use in mainland China. In key hubs such as Xi'an, water shortages driven by rising demand, pollution and upstream withdrawals may necessitate relocation to southern.

REGULATIONS

Regulatory scrutiny of data centre sustainability is increasing globally as governments respond to the sector's rapidly growing energy, water, and land-use footprint. Data centres are no longer treated solely as digital infrastructure but are increasingly regulated as resource intensive industrial assets, subject to environmental permitting, efficiency standards, and disclosure requirements. Policymakers are focusing in particular on energy efficiency, water use, waste heat recovery, emissions, and resilience, reflecting concerns around grid capacity, water stress, and climate targets.

As a result, regulation is evolving from largely voluntary best practice frameworks toward mandatory reporting, performance thresholds, and conditional access to permits, public funding, and procurement. This shift materially raises compliance expectations for developers and operators and is beginning to influence siting decisions, design standards, and operating models across major data centre markets.

¹⁵ TNFD (2026), Nature-related Issues in the Technology Sector. Dependence on water by the semiconductor and data centre industries.

Case Study: Europe¹⁶

Europe has one of the most developed sustainability policy environments affecting data centres, though it remains a mix of binding horizontal legislation and sector-specific obligations that are still evolving. At present, EU-level regulation is anchored primarily in the Energy Efficiency Directive (EED), which requires large data centres to report information on their energy performance and sustainability characteristics. The Directive requires the disclosure of sustainability KPIs, waste heat reuse and its technical and economic feasibility, requirements for energy audits and environmental/energy management systems and general requirements for member states. It also encourages the EU Code of Conduct on Data Centre Energy Efficiency. A draft labelling / rating scheme for data centres is expected to be proposed in April, which may impose Power Usage Effectiveness (PUE) and water consumption targets, sustainability metrics linked to access to financing and it will also likely have an impact on government procurement. More specifically, Germany is looking to implement minimum PUE and renewable energy requirements and Spain has a proposal for additional assessments of data centre's impact on society.

WATER AS A POLITICAL AND COMMUNITY ISSUE

Water in the United States is not just a physical resource but a deep political and community issue because decisions about its management determine who bears the risks of scarcity, pollution, and climate impacts. Nearly 10% of Americans live in areas with limited water supplies, with socially vulnerable communities disproportionately affected, while widespread contamination of aquifers and waterways - by arsenic, manganese, radionuclides, and nitrate - threatens public health, especially for low-income and minority households reliant on domestic wells. These pressures are intensifying as climate change drives extreme heat, drought, flooding, reduced snowpack, and saltwater intrusion, undermining both water quality and availability.¹⁷

At the same time, competing demands shape political choices: agriculture remains the largest water user, consuming over 110 billion gallons per day, even as groundwater in regions such as the High Plains continues to be depleted at alarming rates. Federal policy direction also matters, with recent rollbacks of environmental protections raising concerns about the adequacy of national action on water and climate risks.¹⁸

Importantly, water has been known to drive conflict, including 1934 with Arizona and California went to war over water, with Arizona sending National Guard troops to its border. We also can see the potential for water to be weaponised as with the recent targeting of Middle Eastern desalination plants. This underscores the need to think of robust water supply chains that will last over time.

¹⁶ Covington (2026) - Sustainable AI Infrastructure in the EU

¹⁷ [Water supplies strained for 30 million Americans, new USGS report warns of national water crisis | NationofChange](#) (2025)

¹⁸ [USGS releases a comprehensive look at water resources in the United States | U.S. Geological Survey](#) (2025)

Source: [Thirsty for data: water scarcity challenges Gulf's AI ambitions | AGBI & Vital desalination plant in the Gulf become war targets - France 24](#)

TECHNOLOGY PATHWAYS TO MITIGATE WATER RISK

Within these water-related constraint lies a powerful opportunity. History shows that periods of rapid industrial change have repeatedly forced societies to confront the limits of their water systems - and, in doing so, to modernise them. From the emergence of municipal wastewater treatment to the build-out of large-scale water and energy infrastructure, past industrial revolutions triggered “water transitions” that ultimately supported both economic growth and human development. The AI era is now creating a similar inflection point.

Managed poorly, rising AI-driven water demand risks becoming a zero-sum contest between industry and communities, exacerbating local water stress and social opposition. Managed well, however, it can act as a catalyst for a new phase of water innovation - accelerating investment in efficiency, reuse, digital monitoring, and collaborative water management. In this sense, the AI economy does not merely depend on water security; it can help drive it. The need to rethink how water is sourced, treated, reused, and governed for the AI era creates a rare opportunity to build systems that are more resilient, more circular, and better aligned with long-term economic and environmental sustainability.

Water recycling and reuse in semiconductor fabs

Current spending on water infrastructure

Approximately 68% of water infrastructure investment is directed toward water transport, with 39% allocated to pipes and 23% to pumps. Water treatment comprises 17% of expenditure, encompassing a variety of methods from basic to advanced processes, including filtration, organic treatments, and desalination. Metering represents just 2% of water infrastructure spending, yet is growing and is essential for effective and sustainable water management. The remaining 14% covers a broad spectrum of products and services related to the treatment and transmission of water and wastewater.

Water recycling is increasingly central to the resilience of semiconductor fabs. Empirical evidence from a detailed fab case study by Lu et al. (2018) shows how advanced water network design and regeneration can materially reduce this burden. Using a multi-constraint linear programming model applied to a high-intensity microchip fab, the study found that installing additional regeneration units increased the total volume of reused water by 40.1%, raising fab-level water reuse efficiency from 83.5% to 87.2% while cutting discharge rates from 62.9% to 57.5%.¹⁹ AI will likely help find new efficiencies.

The authors note that industry targets already call for water recycling and reclamation rates above 75%, reflecting both regulatory pressure and business risk management in drought-prone regions. Crucially, the study also demonstrates that higher water tariffs and tighter discharge limits significantly strengthen the economic case for recycling, as reclaimed water can be cheaper than sourcing and disposing of freshwater. Water recycling in fabs is no longer incremental efficiency improvement, but a system-level strategy that safeguards production continuity, reduces exposure to climate-driven water scarcity, and under the right conditions, lowers long-term operating costs.

19 [Strategic optimization of water reuse in wafer fabs via multi-constraint linear programming technique - ScienceDirect](#) (2018)

Source: Barclays (2026). Semis' thirst for water: A shifting global picture

Liquid cooling and closed loop systems in data centres

The principle of avoidance first, reuse through recirculation second, is now reshaping cooling strategies in high-density data centres, where rising AI workloads are pushing conventional air cooling to its physical limits. Traditional evaporative cooling systems can consume up to 1.5 million litres of water per day, directly exposing operators to water scarcity risks.²⁰ In response, companies such as Microsoft, Vertiv, Evolution Data Centres, and Bridge Data Centres are deploying closed-loop liquid and hybrid cooling systems that eliminate ongoing water use altogether.²¹ Vertiv's hybrid X-Cooling systems and immersion-cooling technologies pioneered by firms such as Iceotope remove heat without evaporation, pumps, or freshwater withdrawals, while delivering major efficiency gains for AI workloads.

This operational shift is reinforced by recent thermal-engineering evidence showing that air cooling is approaching its feasible performance limit, while closed-loop liquid systems can reliably support far higher power densities (of 75kW+). In controlled experiments, Naduvilakath-Mohammed et al. (2023) demonstrate that closed-loop liquid cooling can maintain safe chip temperatures at power roughly three times the practical limit of fan-based air cooling - while remaining energy-efficient when fluid flow and fan speeds are properly optimised.²²

Digital water management and leakage reduction

AI and smart metering can play a powerful role in reducing water losses by enabling real-time monitoring, predictive maintenance and early leak detection across increasingly complex water systems: by analysing high-frequency data from meters, pressure sensors, satellites and operational records, AI can identify anomalies that signal emerging leaks or infrastructure failures, optimise pumping schedules and pressure management, and help utilities target maintenance and capital investment more efficiently, reducing both water waste and energy costs. However, this creates a paradox: the same AI technologies that improve water efficiency are themselves driving a growing water footprint, largely through the water-intensive data centres required to train and operate advanced AI models, as well as the indirect water demands associated with their rapidly rising energy consumption and material supply chains. As AI adoption accelerates, communities may therefore experience simultaneous gains from smarter, more efficient water management and new pressures on local water resources from AI-enabled industrial growth, underscoring the need to align digital water solutions with low-water data centre design, water reuse, and closer collaboration between utilities, municipalities and large technology users.²³

20 [How are Data Centres Shifting to Zero-Water Cooling Tech? | Data Centre Magazine](#) (2025)

21 [How are Data Centres Shifting to Zero-Water Cooling Tech? | Data Centre Magazine](#) (2025)

22 [Closed loop liquid cooling of high-powered CPUs: A case study on cooling performance and energy optimization - ScienceDirect](#) (2023)

23 [AI in water management: Balancing innovation and consumption | White & Case LLP](#) (2025)

Limits of technology: what efficiency cannot solve²⁴

Even where water recycling is improving across the AI value chain, this does not eliminate the fundamental impact of large initial water withdrawals for data centre cooling. Unlike household uses such as dishwashers or toilets - where water typically returns to wastewater treatment facilities and is recycled back into the local supply - much of the water used in data-centre cooling systems is consumed, not merely used.

In evaporative cooling systems, water is deliberately converted into water vapour to remove heat and is released into the atmosphere. Once vaporised, this water no longer returns directly to local rivers, aquifers, or reservoirs. While it remains part of the global hydrological cycle and may eventually fall as precipitation, it is effectively lost to the local water balance in the short to medium term. Not all atmospheric moisture returns as recoverable rainfall, and even when it does, it may fall far from the point of extraction or over long time horizons.

This distinction is critical when assessing claims around “recycled” or “reused” water in the AI value chain. Recycling can reduce the rate of withdrawal, but it does not fully offset the initial large extraction required to operate cooling systems at scale. As a result, data-centre water use can contribute to declining regional water tables, particularly in water-stressed or drought-prone areas-where a large share of new data-centre capacity has been built since 2022.

Because the return of this water to local systems is slow and uncertain, cooling-related water use functions as effectively non-renewable in the short term. This local depletion can have knock-on effects, including reduced groundwater availability, heightened competition with other users, and potential impacts on local microclimates and weather patterns. In this sense, even as efficiency and recycling improve, the physical realities of evaporation mean that data centre water use remains.

INVESTMENT IMPLICATIONS

Water risk as a margin, capex and growth constraint

The increasing scarcity of water can lead to several operational constraints for companies that have significant exposure.

For existing operations, there are several ways in which this can act as a headwind on margins for companies as the costs of maintaining existing operations increases. The most obvious example would be that increasing scarcity should drive up utility prices of water over time, adding a direct cost. We saw in 2020-21 for example that Taiwan was faced with water shortages and a severe drought; this caused TSMC to be faced with a 15% restriction in water supply (even after having their water use prioritised over other sectors), and ultimately the company spent USD 28.6mm on water trucks to maintain their production schedule – a 60% YoY increase in their emergency budget for water.

Indirectly, there may be additional incurred costs as a result of needing to comply with increasing water quality and handling standards, such as investing in water treatment and monitoring. Additionally in a worst-case scenario of rationing water supply, a lack of availability of water may reduce the ability for companies to generate the same level of output which would hurt margins as overall utilisation of factories and plants is reduced – especially given that in the asset-heavy semiconductor manufacturing industry, utilisation is a key determinant of margins.

This then also becomes a more difficult question for companies assessing capex or growth opportunities, inherently as the uncertainty around water availability creates question-marks over the return on any incremental investment. For the reasons outlined above, capex

²⁴ [AI is gobbling up water it cannot replace – I’m working on a solution](#) (2025)

on future projects may be higher given the need to invest in more systems and processes around water management. From an execution perspective, projects may also take longer to materialise as they may need to go through more steps on planning and permitting that were previously not as significant a consideration. These arguments would in theory reduce returns on future capex.

In some areas where there are significant water constraints, the lack of available incremental resource may simply create a limit on how much existing companies can expand, or even over time create a stranded asset risk for companies, as plants can no longer continue to operate in water-scarce environments.

These risks are also likely to be compounded over time for operators of AI data centres, given that the increasing power of semiconductors used for LLM Training is likely to require a transition to direct-to-chip liquid cooling solutions, increasing the amount of water required to operate a data centre. An estimate from Xylem and Global Water Intelligence reported that annual water withdrawals for data centres could more than triple from 2030 to 2050 as a result of more AI applications.

Finally, this does not consider the realistic impact that water scarcity would have on prices and inflation for consumers- it is fair to assume that companies facing margin pressure from rising commodity prices would aim to negotiate with their customers to sell at higher prices, ultimately preserving some of their margin. As a result, water scarcity could ultimately impact end-consumers through higher prices of electronics that rely on chips that are now more expensive to manufacture, or services provided by data centres that are seeing increased costs due to water usage and cooling requirements.

Who is most exposed along the value chain?

Companies with the most risk in the semiconductor industry would have a combination of high demands of water for manufacturing processes, located in high-stress areas for water, as well as with limited flexibility to relocate (for example, due to supply chain constraints, especially in sectors such as semiconductors where the supply chain can be quite extended with various components and services).

Research analysing the risk across semiconductor manufacturing facilities from water found that in a 'Business As Usual' climate scenario, by 2030 at least 40% of facilities globally are in areas of extremely high or high water stress risk.

Barclays Research has estimated that of the major semiconductor and memory manufacturers, SK Hynix, ST Micro, and Samsung Electronics all have >10% of their water withdrawals coming from water-stressed regions. In a high-impact scenario, Barclays also estimates that increased stress could reduce Infineon's and ST Micro's annual net income by 4% and 7% respectively (on 2027 estimates).

For data centre operators / hyperscalers, aside from the inherent issues around the supply and cost of semiconductors as indicated above, we can assess the exposure of data centre operators to areas of high water stress. Bloomberg intelligence found that Tencent and Iron Mountain have the greatest exposure with >80% of their total capacity in high-stress regions. IREN and China Unicom also have high exposure. Larger US hyperscalers like Meta and Microsoft are relatively less exposed but still have c.40% of their capacity in high-stress areas, although IBM sits at roughly 90%

This is largely going to remain an issue for hyperscalers; further research indicates that for several companies c.50% of data centre capacity that is in construction area in areas of high or extreme water stress.

Beneficiaries: water treatment, recycling, cooling technologies

The beneficiaries of this trend will be involved in the supply chain of water that can benefit from the growing demand for water use in the technology sector. However, other companies can more broadly benefit from this by working to design more efficient systems around the data centre that reduces the demand intensity of water.

We identify key categories of companies which should benefit:

- **Specialised water solutions** such as companies that provide ultrapure water for the semiconductor manufacturing process, that are crucial for manufacturing semiconductors and will also be increasingly needed as the complexity of manufacturing leading-edge semiconductors requires more water; specific companies include Organo, Kurita Water, Nomura Micro Science.
- **Circuit water solutions** (recycling, reuse, treatment): this could for example be companies such as Ecolab, who through their Nalco unit offers water management solutions to data centres, or Veolia who offers water treatments like reverse osmosis systems and deionization systems to customer in this sector. This would also include components companies for these solutions such as Toray Industries which produces specialised membranes in water treatment systems.
- **Water cooling solutions** who can work to optimise water use within the data centre. Air cooling / traditional cooling exposed stocks include names in the HVAC space such as Johnson Controls and Munters. Key companies in the liquid cooling value chain include e.g. Vertiv, NVent, and Delta Electronics, but also more recently companies such as Eaton and Schneider who have acquired key private companies in the space such as Boyd and Motivair.
- **Water efficiency solutions:** companies which can find innovations elsewhere in the data centre or tech landscape that would reduce the need for water use. A recent example would be the announcement from Nvidia in early 2026 where they announced that the next generation of their AI chips, Vera Rubin, would be compatible with cooling at higher water temperatures of 45 degrees Celsius.

Differentiating leaders from laggards

Differentiating leaders from laggards in the above bullet points can be challenging given that the semiconductor and AI data centre ecosystem is rapidly evolving, as a result of a faster cadence of development at the AI chip level than in traditional semiconductors.

While there is not a universal path for success, leading companies may feature:

- Strong knowledge around a niche, critical area that cannot be substituted in the supply chain, or companies with strong R&D teams that can develop cutting edge products that can improve water use in the future;
- A track record of developing water-related technologies and optimising recycling/reuse systems;
- A global presence to provide a greater reach to customers that have physical assets located around the world in different environments;
- A broad product portfolio that can span several parts of the water and cooling value chain, to find efficiencies and optimisations at the system-level as well as at the product level;
- Close cooperation with the market leaders in the semiconductor industry to gain an insight into what the future roadmap for water use looks like and working with these companies to collaboratively design the systems that will be adopted in the future.

CASE STUDY - SUCCESSFUL MITIGATION STRATEGIES IN PRACTICE

Organo Corp – Integrated Water Solutions for Semiconductors and Data Centres

Organo addresses the challenge of water efficiency, purity, recycling, and operational control by delivering integrated water treatment, recycling, and control solutions tailored to semiconductor fabs and data centre operations.

Semiconductor Manufacturing Water Solutions

Organo provides end-to-end water infrastructure for semiconductor manufacturing, including ultrapure water (UPW) supply systems for wafer and chip cleaning, advanced wastewater treatment facilities, and closed-loop water recycling technologies. Their solutions enable more than 80% of water to be recovered and reused, significantly reducing freshwater withdrawals. Beyond water recycling, Organo's systems also recover valuable materials from wastewater, transforming waste streams into resource streams and improving overall process efficiency. These capabilities are particularly relevant for the electronics and semiconductor sectors, where demand is accelerating in response to the growth of AI and high-performance computing.

Data Centre Cooling Solutions

In parallel, Organo Corp supports water and energy efficiency in data centres and industrial facilities through cooling water treatment solutions designed to reduce both energy use and environmental impact. The company offers high-performance treatment chemicals with strong bactericidal effects, alongside automatic control and remote monitoring technologies that optimize system performance. A flagship example is ORCHASER V, an automatic control device that precisely manages chemical concentration in cooling water systems. By maintaining optimal dosing levels, ORCHASER V helps lower chemical consumption, reduce maintenance requirements, and minimize environmental impact while improving operational reliability.

WHAT INVESTORS SHOULD WATCH NEXT

Leading indicators of water stress

One of the leading indicators of water stress, is to assess whether water-related issues are creating environmental or permitting delays on current projects due to a lack of available resource. Equally, monitoring existing assets to see if they are putting additional strain on the availability of water for other sectors or households in local areas may be worth monitoring, as this may lead to a broader discussion around these companies' 'License to Operate' – especially if they are crowding out consumer water demand.

Over time it will be important to also note if capex spending on AI starts to moderate due to water constraint issues – however this may be unlikely in the shorter term given other bottlenecks currently constraining data centres more than water (e.g., availability of chips, power generation, and labour).

Water metrics that matter (and those that don't)

Water Withdrawal: defined as water removed from the ground or diverted from a surface-water source for use. This is obviously a key metric to track how impactful a company's water demands are on the wider environment.

Water Recycling Rate: is a useful metric to track that companies may report and target. This can be used alongside the water withdrawals data to assess how much progress a company could make with respect to handling their water demands.

Water Usage Effectiveness (WUE) – more applicable for data centre operators, measures the total water consumed by a data centre, with respect to the energy consumed by the IT equipment (i.e. chips, servers). This tells us how efficiently data centres are by scaling water use to compute requirements. Typical units are e.g. Liters/KWh. According to the US Department of Energy, in the US The average WUE is 1.8 litres/KWh.

WUE is usually used in combination with other metrics such as Power Usage Effectiveness (PUE) and Carbon Usage Effectiveness (CUE) to holistically assess the overall resource efficiency of a data centre.

We find WUE a preferable metric to assess the intensity of water use as the energy consumption of a data centre is a more fundamental operational metric tied to the same physical assets. In contrast, using a financial metric such as Net Income or Revenues to scale a water metric may create difficulties in interpretation given the multitude of factors that impact such financial metrics (e.g., industry price trends, operational leverage, financial leverage and tax differences between companies, impact from unrelated business units within the company, etc).

For semiconductor manufacturers, TSMC reports a similarly-scaled metric of Unit Water Consumption, which in this operational context is measured as litres of water consumption per wafer mask layer produced (the various processes requiring water outlined earlier in the report such as lithography, deposition and etching ultimately are used to produce wafers, but over time with increased complexity we are seeing higher layer counts in each wafer, leading to even more water use).

Risk warning

Investments are subject to market fluctuations and other risks inherent to investing in securities. The value of investments and the income they generate may rise or fall and it is possible that investors may not recover their initial investment

CONCLUSION

Water is becoming a binding constraint on the digital economy as AI-driven growth accelerates demand for data centres and advanced semiconductor manufacturing. This paper shows that water dependence is deeply embedded across both on-site cooling and upstream chip production, with a growing share of critical infrastructure located in regions already facing medium to extreme water stress. While efficiency gains, recycling, and new cooling technologies are reducing water intensity, they do not eliminate the impact of large absolute withdrawals - particularly where evaporative cooling leads to local depletion that is not quickly reversible. As a result, water risk is shifting from a peripheral environmental consideration to a material factor influencing operating costs, permitting, capital allocation, and the feasibility of future expansion.

At the same time, this constraint creates opportunity. Rising water pressure is accelerating investment in circular water systems, advanced treatment, liquid cooling, and digital water management, benefiting companies that enable more resilient and efficient infrastructure across the value chain. For investors, incorporating water risk into analysis is increasingly essential to distinguishing leaders from laggards and identifying durable sources of return. In the AI era, water security is not only a sustainability issue - it is a prerequisite for long-term growth, resilience, and value creation in the digital economy.



BNP PARIBAS ASSET MANAGEMENT Europe, "the investment management company", is a simplified joint stock company with its registered office at 1 boulevard Haussmann 75009 Paris, France, RCS Paris 319 378 832, registered with the "Autorité des marchés financiers" under number GP 96002. This material is issued and has been prepared by the investment management company. This material is produced for information purposes only and does not constitute: 1. an offer to buy nor a solicitation to sell, nor shall it form the basis of or be relied upon in connection with any contract or commitment whatsoever or 2. investment advice. Opinions included in this material constitute the judgement of the investment management company at the time specified and may be subject to change without notice. The investment management company is not obliged to update or alter the information or opinions contained within this material. Investors should consult their own legal and tax advisors in respect of legal, accounting, domicile and tax advice prior to investing in the financial instrument(s) in order to make an independent determination of the suitability and consequences of an investment therein, if permitted. Please note that different types of investments, if contained within this material, involve varying degrees of risk and there can be no assurance that any specific investment may either be suitable, appropriate or profitable for an investor's investment portfolio. Given the economic and market risks, there can be no assurance that the financial instrument(s) will achieve its/their investment objectives. Returns may be affected by, amongst other things, investment strategies or objectives of the financial instrument(s) and material market and economic conditions, including interest rates, market terms and general market conditions. The different strategies applied to the financial instruments may have a significant effect on the results portrayed in this material. All information referred to in the present document is available on www.bnpparibas-am

April 2026 - Design: CREATIVE SERVICES BNPP-AM - P2604003

VIEWPOINT



BNP PARIBAS
ASSET MANAGEMENT

**The sustainable
investor for a
changing world**